

Poverty Alleviation Through Geographic Targeting: How Much Does Disaggregation Help?

Chris Elbers, Tomoki Fujii, Peter Lanjouw, Berk Özler and Wesley Yin¹

August 18, 2004

In this paper, using recently completed “poverty maps” for three countries, we simulate the impact on poverty of transferring an exogenously given budget to geographically defined subgroups of the population according to their relative poverty status. We find large gains from targeting smaller administrative units, such as districts or villages. However, these gains are still far from the poverty reduction that would be possible had the planners had access to information on household level income or consumption. Our results suggest that a useful way forward might be to combine fine geographic targeting using a poverty map with within-community targeting mechanisms.

Key Words: Targeting, Poverty, Poverty Maps.

JEL Classification Numbers: C15, I32, H53.

¹ Elbers is with Free University, Amsterdam, Fujii is with University of California, Berkeley, Lanjouw and Özler are with the World Bank, and Yin is with Princeton University. The authors are grateful to Martin Ravallion, Norbert Schady, and participants in the “Poverty and Applied Microeconomics” seminar series at the World Bank for useful discussions and comments. These are the views of the authors and should not be attributed to the World Bank or any affiliated organization. Address for correspondence: planjouw@worldbank.org and bozler@worldbank.org.

I. Introduction

Public policies in developing countries are often articulated in terms of poverty reduction objectives. Resources for such purposes are invariably scarce relative to the number and magnitude of competing claims. Spending priorities must be defined, and it is often desirable to target social transfers to those beneficiaries whose needs are most urgent. Coady and Morley (2003) survey experience with such targeted transfer programs and show that errors of inclusion and exclusion are unavoidable consequences of such targeting efforts. Efforts aimed at improving targeting of public spending generally focus on reducing either one, or sometimes both, of these types of errors.

Because the precise economic circumstances of households can be difficult to ascertain it is not easy to define who should be eligible to receive a government transfer. Nor is it straightforward to design an administrative mechanism to ensure that the transfer actually reaches the intended beneficiary. In practice governments often exploit geographic variability in the design of targeting schemes: poverty is typically thought to be more concentrated in some areas of a country than others and most countries have an administrative structure that disaggregates to different levels. For example, the central government, located in the capital city, may rely on state or provincial governments to implement government policies at the state or province-level. These administrations might rely, in turn, on counties or districts, which may themselves rely on yet lower levels of administration. Resources aimed at poverty reduction can thus be directed to those localities where poverty is concentrated and administration of these transfer schemes can be carried out at the relevant local level.

Despite their intuitive appeal, transfer schemes that target poor communities remain difficult to design. One of the difficulties concerns information on the spatial distribution of poverty. In many developing countries there is now a good deal of experience with methods to conceptualize, measure, and analyze poverty. Yet most of the empirical evidence that has been brought to bear on this topic tends to be at the national level. This is primarily due to the fact that data on incomes or consumption expenditures, which serve as prime input into the quantitative analysis of poverty, tend to derive from sample surveys of households. The sample size of such surveys is often rather small. This implies that statements about the spatial distribution of poverty can generally refer only to very broad geographic breakdowns. It is not uncommon, for example, for surveys to be “representative” only for urban and rural areas of a country, in combination with perhaps a very crude geographical breakdown of the country into broad regions (north/south or hills/plains). While there is a general suspicion in many countries that poverty occurs in geographic “pockets” that are defined fairly precisely, sample surveys do not readily allow one to confirm or refute such notions.²

Absent detailed information on poverty outcomes at the local level, policymakers interested in targeting transfers to the local level have often sought to use proxy indicators of local poverty outcomes. For example, the Microregions program that is being contemplated by SEDESOL (Ministry of Social Development) in Mexico is proposing to deliver social services and to promote productive investment projects in a set of municipalities that have been selected on the basis of a marginality index produced at the municipality level by CONAPO (de Janvry, Sadoulet, Calva and Solaga, 2003). This index is based on a set of simple variables taken from the population census, and is believed to be correlated with conventional notions of poverty.

² However, recent research shows that beliefs regarding the existence of such “pockets of homogeneous poverty” may not always be true. Elbers et al. (2004) demonstrate that small communities in three developing countries are found to vary markedly from one

When such proxy indicators are used for targeting instead of explicit poverty estimates, there is mis-targeting both due to the targeting errors described above and due to problems with the proxy welfare index at the community level.³

In recent years there has been growing experience with the development of “poverty maps” that explicitly estimate consumption or income-based welfare outcomes at the local level.⁴ This approach involves imputing consumption or income at the unit-record level into the population census based on a set of regression models estimated from household survey data. The imputed household-level income or consumption estimates are then aggregated into welfare indices (including but not confined to poverty measures) at different levels of geographic aggregation. While there is inevitably a degree of statistical error associated with these welfare indices, this uncertainty can be quantified, and it is not uncommon to find that the method produces reliable estimates of poverty for communities comprised of only 1,000-5,000 households on average.⁵

This paper asks to what extent the high degree of spatial disaggregation offered by such poverty maps in three different countries can help to improve targeting schemes aimed at reducing poverty. In this sense, the paper closely builds on the earlier analysis in Ravallion (1993) who finds that spatial disaggregation to the broad regional level in Indonesia – the lowest level at which household survey data provide reliable estimates of poverty – improves targeting but only to a modest extent. As in Ravallion (1993), we consider the distribution of a hypothetical budget to a country’s population, assuming that we have no information about the poverty status of this population other than the geographic location of residence and the level

another in terms of the degree of inequality they exhibit and that there should be no presumption that inequality is less severe in poor communities.

³ For an example of the latter, see Hentschel et al. (2000).

⁴ Hentschel et al, 2000, and Elbers, Lanjouw and Lanjouw, 2002, 2003 describe the basic methodology. At present, the methodology has been implemented or is in the process of being implemented in some 30 developing countries. For further details see the website:

of poverty in each location. As a benchmark case with which to compare our results, we make the extreme assumption of no knowledge whatsoever about the spatial distribution of poverty – in which case our given budget is distributed uniformly to the entire population. We set up a series of comparisons to this benchmark, where we assume knowledge about poverty levels for administrative units with progressively smaller populations. For a given level of disaggregation, we ask how knowledge about poverty outcomes across regions can be incorporated into the design of a transfer scheme so as to improve the overall targeting performance relative to the benchmark case. We consider a variety of transfer schemes that make use of this knowledge in different ways. The schemes range from simple, intuitive transfer schemes to more sophisticated ones, in the latter of which expected poverty at the national level is minimized given the information and budget constraints. We are interested in comparing performance across schemes, as well as the relative performance of each scheme at alternative levels of disaggregation. We consider performance in terms of the squared poverty gap – a measure of poverty that increases not only with the number of people below the poverty line but is also particularly sensitive to the distance between a poor person’s income level and the poverty line. To avoid making strong assumptions about the exact value of the poverty line or the size of the budget, we also show results for two poverty lines and two hypothetical budgets.⁶ Finally, we report how close “optimal geographic targeting” in combination with our poverty maps, comes in terms of poverty reduction to the hypothetical scenario of “perfect targeting”. From this we can get a sense of the potential benefit from combining detailed geographic targeting with additional targeting mechanisms such as

<http://econ.worldbank.org/programs/2473/topic/14460/>.

⁵ For a more detailed discussion, see Demombynes et al. (2002).

⁶ We have tried more poverty lines and budget sizes, but do not present them in this paper for brevity. For a formal discussion of using “program dominance curves” to assess the poverty impact of different programs, see Duclos, Makdissi, and Wodon (2003).

individual or household level means-testing or the incorporation of self-selection targeting mechanisms within communities.

Our simulations are carried out using recently produced poverty maps for Ecuador, Madagascar and Cambodia. These countries are obviously highly heterogeneous in terms of their geographic location, but have also very different social and political structures and are at different stages of overall development. We are interested to examine whether there exist commonalities across these countries in terms of the degree to which the availability of poverty data at the local level can contribute to improvements in the targeting of public resources.

We find that there are potentially large gains to targeting performance from disaggregating to the local level. The benefits become increasingly evident as one makes use of more and more disaggregated data on poverty. In all three countries examined, we show that relative to a uniform transfer the same impact on poverty can be achieved at considerably less expense when targeting is based on the highly disaggregated poverty estimates that are available in poverty maps. The gains are generally more muted when the targeting scheme makes only crude use of the local level poverty estimates. They are also lower if the poverty line in a given country is particularly high, or if the budget available for transfer is particularly large. In all countries we find as well that despite the gains from geographic targeting, our inability to target households directly implies that overall targeting performance remains far from perfect. This implies that there may be scope for combining geographic targeting with other targeting methods in order to reduce errors of inclusion and exclusion even further.

The results in this paper are likely to be of use to policy makers interested in the design of transfer schemes aimed at reducing poverty. It is important to emphasize, however, that there are important caveats that attach to these results. First, we assume that the willingness of

government to consider geographic targeting implies a willingness to sacrifice *horizontal equality* in favor of improved targeting efficiency. In other words, the government is willing to accept that households with equal pre-transfer per-capita consumption levels might enjoy different post-transfer consumption levels. Second, we assume in this paper that the budget available for distribution is exogenously determined. We abstract away entirely from the question of how the transfers are to be financed. As has been argued by Gelbach and Pritchett (2002), political economy considerations are likely to influence options for resource mobilization. It is possible, for example, that fine geographic targeting of transfers is less appealing to voters than a uniform transfer scheme – and this could translate into lower overall budgets available for such targeted transfers than would be the case if the transfer scheme were uniform. Third, we do not address the very real possibility that the costs of administering a given transfer scheme may increase with the degree of disaggregation. It is quite possible that each unit of administration incurs some fixed costs in terms of staff, equipment, and so on, in order to implement a given transfer scheme. Relying on ever lower levels of government to administer the transfer scheme could raise overall administration costs, thereby reducing the overall amount left to transfer. Fourth, it is well recognized that the availability of transfers to certain groups in the population may induce behavioral responses in the population. Those considered ineligible for the transfer might behave in such a way as to appear eligible. Households might, for example, move to those locations where it is announced that transfers will be directed or they can pretend that they are residents of those locations while living elsewhere. It is possible that the cost, to households, of such migration or cheating is lower when it involves small neighboring communities as opposed to large neighboring regions.

Finally, we have not examined optimal targeting in the context of local political-economy considerations or in the presence of community specific public goods as a result of voluntary contributions at the local level. To the extent that the inequality level in a given locality represents some kind of political economy “equilibrium”, it is not clear how a government transfer to a community will actually get distributed across the population in the community. The inequalities in power and influence that prevail in the community are likely to influence how such transfers are allocated. Such factors are likely to result an overestimation of the impact of the targeting scheme on poverty reduction.⁷ In addition, Kanbur (2003) shows that basic results of the targeting literature can change in the presence of community-specific public goods, and that optimal targeting for poverty alleviation can lead to paradoxical results for certain values of the poverty aversion parameter, for example that targeting transfers to the richer community can result in greater welfare gains for the poor (via the increased provision of public goods by the richer segments).

As a result of these caveats, it is important to stress that the gains from fine geographic targeting illustrated in this paper should be viewed as illustrative only. These *potential* gains should be juxtaposed against the potential costs of such targeting and political-economy considerations. Policymakers need to assess such programs on a case-by-case basis to determine whether fine-geographic targeting is the appropriate strategy.

Our paper is structured as follows. In the next section we briefly summarize the methodology and data underpinning the poverty map estimates in Ecuador, Madagascar and Cambodia. We emphasize that the spatially disaggregated poverty data available to us are estimates, with confidence bounds, rather than actual measures of poverty. We indicate how

⁷ On the other hand, it is also possible that the infusion of transfers into a poor community would increase risk-sharing in that community and thereby contribute to further reductions in poverty.

we incorporate this imprecision into our simulation analysis. Section III describes the different targeting schemes that are assessed in the simulation stage. In this section we also demonstrate how one particular targeting scheme can be viewed as optimal in terms of ensuring the maximum possible gains from geographic targeting. Section IV presents the results from our simulation analysis, and Section V presents a concluding discussion.

II. Producing Local Estimates of Poverty

The methodology we implement here has been described in detail in Elbers, Lanjouw and Lanjouw (2002, 2003). We estimate poverty based on a household per-capita measure of consumption expenditure, y_h . A model of y_h is estimated using household survey data, restricting explanatory variables to those that are also found in, and strictly comparable to, the population census. The regression models consumption on a set of household-level demographic, occupational and educational variables as well as census variables calculated at the level of the census-tract or other level aggregation above the household level.⁸

Letting W represent an indicator of poverty or inequality, we estimate the expected level of W given the observable characteristics in the population census and parameter estimates from model estimated on the household survey data.

We model the observed log per-capita expenditure for household h as:

$$(1) \quad \ln y_h = \mathbf{x}_h \boldsymbol{\beta} + u_h,$$

where $\mathbf{x}_h \boldsymbol{\beta}$ is a vector of k parameters and u_h is a disturbance term satisfying $E[u_h | x_h] = 0$. The model in (1) is estimated using the survey data. We use these estimates to calculate the welfare of an area or group in the population census. We refer to our target population as a ‘region’.

Because the disturbances for households in the target population are always unknown, we consider estimating the expected value of the indicator given the census households' observable characteristics and the model of expenditure in (1). We denote this expectation as

$$(2) \quad \mu_v = E[W | \mathbf{X}_v, \boldsymbol{\xi}],$$

where $\boldsymbol{\xi}$ is the vector of model parameters, including those that describe the distribution of the disturbances.

In constructing an estimator of μ_v we replace the unknown vector $\boldsymbol{\xi}$ with consistent estimators, $\hat{\boldsymbol{\xi}}$, from the survey-based consumption regression. This yields $\hat{\mu}_v$. This expectation is generally analytically intractable so we use simulation to obtain our estimator, $\tilde{\mu}_v$.

The difference between $\tilde{\mu}_v$, our estimator of the expected value of W for the region, and the actual level of welfare for the region reflects three components. The first, (*idiosyncratic error*), is due to the presence of a disturbance term in the first stage model which implies that households' actual expenditures deviate from their expected values. This component becomes increasingly important as the target population becomes very small. The second component of our prediction error is due to variance in the first-stage estimates of the parameters of the expenditure model (*model error*). We calculate the variance due to model error using the delta method (see Elbers et al 2002, 2003). The third component of our prediction error is due to inexact method to compute $\hat{\mu}$ (*computation error*). This component can be set arbitrarily small by choosing a large enough set of simulation draws.

⁸ In the case of Madagascar and Cambodia, we also include regressors from tertiary datasets in the regression model (see Mistiaen, Ozler, Razafimanantena, Razafindravonona, 2002, and Fujii, 2003),

Implementation

The first-stage estimation is carried out using household survey data in our three respective countries. These surveys are stratified at the region or state level, as well as for rural and urban areas. Within each region there are further levels of stratification, and also clustering. At the final level, a small number of households (a cluster) are randomly selected from a census enumeration area.

Our empirical model of household consumption allows for an intra-cluster correlation in the disturbances (see Elbers, Lanjouw and Lanjouw, 2002, 2003 for more details). Failing to take account of spatial correlation in the disturbances would result in underestimated standard errors. We estimate different models for each region and we include in our specification census mean variables and other aggregate level variables cluster-level effects. All regressions are estimated with household weights. We also model heteroskedasticity in the household-specific part of the residual, limiting the number of explanatory variables to be cautious about overfitting. We approximate both the cluster- and household-level disturbances as either normal or t distributions with varying degrees of freedom.⁹ Before proceeding to simulation, the estimated variance-covariance matrix is used to obtain GLS estimates of the first-stage parameters and their variance.

Data

The data used in this study consists of a household survey and a population census from each of Ecuador, Madagascar, and Cambodia. Table 1 presents the basics on each of the data

⁹ Rather than drawing from parametric distributions in our simulations, we can also employ a semi-parametric approach by drawing from observed residuals in the first stage model. Our results have generally been found to be quite robust to the choice of parametric or semi-parametric draws.

sources, such as year, sample size, stratification, etc. For more detail on the data, please refer to the studies listed in the “References” row in Table 1.

III. Transfer Schemes and Simulation Procedures

The transfer schemes

As described in Section I, we compare a variety of targeting schemes against a benchmark scheme which assumes absolutely no knowledge of the geographic distribution of poverty. One of our main objectives is to see whether, and to what extent the availability of poverty estimates for different geographic locations can help to improve the poverty impact of distributing a given budget. In our baseline, benchmark, case we postulate that the government has a budget, S , available for distribution and wishes to transfer this budget in such a way as to reduce poverty. However, because the government is assumed to have no knowledge of who the poor are or where they are located, it is unable to distribute its budget in any manner other than a lump-sum transfer to the entire population of size N . In our benchmark case, thus, we calculate the impact of transferring S/N to the entire population.¹⁰

Optimal use of geographically disaggregated information on poverty has been investigated by Kanbur (1987), Ravallion and Chao (1988), Glewwe (1991), Ravallion (1993), and Baker & Grosh (1994). Kanbur (1987) formalized the theoretical problem of policy design under imperfect information, while Ravallion & Chao (1998) demonstrated how this general targeting problem can be solved in a computationally feasible way.¹¹ For our purposes, the

¹⁰ It could be argued that our benchmark scenario is not terribly realistic. Perhaps more likely would be a situation where absence of detailed information on the extent and distribution of poverty, and absent any specific effort to target the poor, would result in a default situation of resources being appropriated by the non-poor (see the discussion in Campante and Ferreira, 2003). To the extent that this is true our estimates of the gains from targeting, once we assume some information on the distribution of poverty, might be seen as conservative estimates of the true benefits.

¹¹ As we use predicted expenditures from census data unlike Ravallion and Chao (1988), who use observed income data from household surveys, we utilize a different algorithm to solve the optimization problem. Applying their algorithm to our setting would yield the same results.

important result from Kanbur (1987) is that if decision makers wish to transfer resources in such a way as to minimize poverty summarized by the Foster-Greer-Thorbecke (FGT) class of poverty measures with parameter value $\alpha > 1$, then on the margin the group with the higher $FGT^{\alpha-1}$ should be targeted. In other words, if the government wishes to minimize the squared poverty gap (equal to a poverty measure from the FGT class with $\alpha=2$), then geographic regions should be ranked by the poverty gap (FGT with $\alpha=1$) and lump-sum transfers made until the poverty gap of the poorest region becomes equal to that in the next poorest region, and then transfers to these two regions should be continued until their poverty gap is equal to the next poorest region, and so on, until the budget is exhausted.

Let y_{ch} denote the per capita expenditure of household h (with m members) living in group c . Assume that the government is able to provide lump-sum supports a_c that differ across groups c . Thus, after-support expenditures are $y_{ch} + a_c$. Suppose the government wishes to minimize expected FGT2 after transfers subject to the constraint that total transfers are limited by the budget S :

$$\min_{a_c} = \sum_c \sum_h m_{ch} \text{Avg}[(y_{ch}^{(r)} + a_c - z)^2 1(y_{ch}^{(r)} + a_c \leq z)]$$

subject to

$$\sum_c a_c \sum_h m_{ch} \leq S$$

$$a_c \geq 0$$

The r^{th} simulated expenditure of household h in c is indicated by $y_{ch}^{(r)}$. The (linear) $\text{Avg}[\]$ operator indicates taking the average over simulations $r=1, \dots, R$ (see below). $1(x)$ is an indicator function, $1(\text{TRUE})=1$, $1(\text{FALSE})=0$. Standard applications of the Kuhn-Tucker conditions leads to first-order conditions for optimal a_c :

$$\sum_h m_{ch} \text{Avg}[(z - y_{ch}^{(r)} - a_c) \mathbb{1}(y_{ch}^{(r)} + a_c \leq z)] \leq \lambda \sum_h m_{ch}$$

$$a_c \left(\lambda - \frac{\sum_h m_{ch} \text{Avg}[(z - y_{ch}^{(r)} - a_c) \mathbb{1}(y_{ch}^{(r)} + a_c \leq z)]}{\sum_h m_{ch}} \right) = 0$$

along with the earlier constraints on a_c . Note that

$$\frac{\sum_h m_{ch} \text{Avg}[(z - y_{ch}^{(r)} - a_c) \mathbb{1}(y_{ch}^{(r)} + a_c \leq z)]}{\sum_h m_{ch}}$$

is nothing but the average simulated after-transfer FGT1 for group c . In other words, in the optimum only the groups with the (after-transfer) highest predicted FGT1 get expenditure support, and those that do receive transfers get the amount which equalizes predicted FGT1.

The second targeting scheme that we compare against our benchmark case assumes some knowledge of the spatial distribution of poverty, but does not make use of this knowledge in any particularly scientific or systematic way. This “naïve” targeting scheme was selected in order to contrast with the “optimal” scheme described above. There are reasons to believe that implementation of an “optimal” scheme will be difficult in practice. It is often important for governments to be able to communicate in a very clear and simple way how resources will be targeted, and this need for transparency and ease of communication may prevent governments from carrying out the fine-tuning needed for an optimal scheme.

Of course, there are virtually an infinity of “naïve” schemes that could be implemented. The scheme implemented here is one, particularly straightforward, example. We have experimented with a variety of alternative, more complicated, versions of this naïve scheme. We have not found any alternative that is obviously more effective. Indeed, the specific

scheme implemented here has the virtue of not only being simple but also surprisingly effective at times.

Our “naïve” scheme takes the following form. We first rank geographic areas by estimated poverty. If our interest was to gauge the impact of our scheme on the headcount rate, we would rank areas by the headcount. But as we wish, in this paper, to assess the impact on the squared poverty gap we rank by those estimates. We have an assessment of overall poverty in the country. We take our budget S and divide it by the total number of poor persons in the country, N_p . Our budget divided by the total number of poor persons yields the transfer a that will be distributed to each person. We select the poorest geographic area and transfer a to all persons in that area. If the budget has not been exhausted in the first region, we move to the next poorest region and transfer a to all persons in this second region. We continue until the budget is exhausted. In the marginal region - that in which the budget is exhausted - we do not transfer a but transfer an equal share of whatever remains in the budget to the population of that last region. Note that this scheme does not guarantee some amount of transfer to all regions. The scheme also implies that households will be receiving differing amounts according to their overall size.¹²

Implementation

Operationalizing our simulation exercises involves taking outputs from the micro-level estimation procedure described in Section II and subjecting these to the simulation procedures described above. To recap, in Section II we described a procedure whereby per-capita

¹² Two alternatives that we have experimented with (but do not report here as results were not noticeably better) include a “naïve poverty-gap” scheme and a “naïve poverty share scheme”. In the former we rank communities by estimated poverty. We then cumulate the poverty gaps amongst the poor within each community. We divide that total by the number of poor people in the region. We then transfer this amount to everyone in the community (poor and non-poor). We carry on until the budget is exhausted, and estimate impact on aggregate poverty. In the latter scheme we calculate community i 's contribution to total poverty. We transfer that

consumption is predicted at the level of each household in the population census, based on a consumption model estimated in the household survey. However, because predicted household-level per capita consumption in the census is a function not only of the parameter estimates from the consumption model estimated in the survey, but also of the precision of these estimates and of those parameters describing the disturbance terms in our consumption model, we do not produce just one predicted consumption level per household in the census. Rather, r predicted expenditures are produced for each household (in our three countries, we have carried out one hundred predictions). For each respective r , parameter estimates are drawn from a multivariate normal distribution that respects the variance-covariance matrices estimated in the survey-based consumption and heteroskedasticity regressions. In addition, disturbance terms at the cluster and household level are drawn from their respective parametric or semi-parametric distributions. These draws are then applied to the census-level regressors and per-capita consumption is predicted. For the next r , a new set of parameters and disturbances are drawn and a new per-capita consumption measure is predicted. The resulting “dump-file” of r predicted expenditures for every single household in the population census is the key database underpinning “poverty maps” and the policy-simulation exercise explored here.¹³

Simulating the impact of uniform targeting

Our baseline, benchmark, policy simulation is calculated in the following way. Budget S is divided by total population N . The resulting transfer a is added to each predicted expenditure

fraction of the budget uniformly to each household in the community. Note that this latter scheme ensures that all communities receive at least something.

¹³ The poverty map estimate of poverty in community, province or region c is produced from this dump file in the following manner: for every replication r , poverty is estimated over all households in c (after weighting by household size). The average of all poverty

in the “dump file”, to yield $y_{ch}^{(r)} + a$. For each replication r we estimate post-transfer national poverty. The average across the r replications of the estimated post-transfer poverty rates yields our expected poverty rate associated with the benchmark, untargeted lump-sum transfer scheme. This new estimated poverty rate can be compared to the original national-level poverty estimate from the poverty map to gauge the impact of the transfer.

Simulating the impact of “optimal” geographic targeting

Simulating the impact of the “optimal” targeting scheme is a bit more complicated. We want to equalize the following expression across the poorest locations of a country:

$$G_c(a_c) = \int_0^z (z - y - a_c)^+ dF_c(y),$$

which is z times the poverty gap in location c , after every person in the location has received a transfer a_c . $F_c(y)$ is the average of the R simulated expenditure distributions of c . The function $(x)^+$ gives the ‘positive part’ of its argument, i.e. $(x)^+ = x$, if x is positive, otherwise 0. Transfers a_c (which must be nonnegative) add up to a given budget S :

$$\sum_c N_c a_c = S,$$

where N_c is the population size of location c . After transfers there is a group of locations all sharing the same (maximum) poverty gap rate in the country. These are the only locations receiving transfers.

We solve this problem by first solving a slightly different problem. Consider the minimum budget $S(G)$ needed to bring down all locations’ poverty gaps to at most the level G/z . This amounts to transferring an amount $a_c(G)$ to locations with before-transfer poverty

estimates, over the r replications, yields the estimated poverty rate in community c , and the standard deviation yields the associated estimated standard error.

gaps above G/z , such that $G_c(a_c(G)) = G$. Once we know how to compute $S(G)$, we simply adjust G until $S(G)$ equals the originally given budget for transfers S . To implement this scheme we must solve the following equation for a_c :

$$G = \int_0^z (z - y - a_c)^+ dF_c(y).$$

In what follows we drop the location index c for ease of notation. Using integration by parts it can be shown that

$$G(a) = \int_0^z (z - y - a)^+ dF(y) = \int_0^{z-a} F(y) dy.$$

In other words we need to compute the surface under the expenditure distribution between expenditure levels $y=0$ and $y=z-t$, for values of t up to z . Instead of computing $G(t)$ exactly, we use a simple approximation. For this to work we split the interval $[0,z]$ in n equal segments and assume that the ‘poverty mapping’ software has generated expected headcounts for poverty lines $z k/n$, where $k=0, \dots, n$. In other words we have a table of $F(z k/n)$. Using the table we approximate $F(y)$ by linear interpolation for y between table values. With the approximated expenditure distribution it is easy to solve for transfers as a function of G (see below). In practice we find that $n=20$ gives sufficiently precise results.¹⁴

The computational set-up is as follows (note that the numbering we adopt means going from z in the direction of 0 rather than the other way around). Define $b_0=0$, and for $k=1, \dots, n$, b_k as the surface under the (approximated) expenditure distribution between $z-kz/n$ and $z-(k-1)z/n$, divided by z :

$$b_k = \frac{1}{2n} (F(z - kz/n) + F(z - (k-1)z/n)).$$

¹⁴ Other interpolation schemes are possible. For instance, if the *poverty gap* is given at table values zk/n an even simpler computation presents itself. Often the poverty mapping software will give percentiles of the expenditure distribution. These can also be used for interpolation, but the formulas are more cumbersome, since the percentiles are not equally spaced.

Let g_0 be the original poverty gap, or in terms of the discussion above, $g_0 = G(0)/z$. For $k=1, \dots, n$, put

$$g_k = g_{k-1} - b_k.$$

The g_k are the poverty gaps of the approximated expenditure distribution for successively lower poverty lines $z - kz/n$. Let a_k be the per capita transfer needed to bring down the poverty line to $z - kz/n$:

$$a_k = kz / n.$$

We can now solve for per capita transfers as a function of the intended poverty gap $g < g_0$:

1. Find k such that $g_{k+1} \leq g < g_k$.
2. The per capita transfers resulting in poverty gap g are

$$a(g) = a_k + \frac{g_k - g}{g_k - g_{k+1}} \cdot \frac{z}{n}.$$

This scheme can be implemented using standard spreadsheet software.

Simulating the impact of “naïve” geographic targeting

Simulating the impact of our “naïve” transfer scheme on the basis of the “dump-file” described above is far more straightforward. We take our poverty map as the basic statement on the distribution of poverty in the country. On the basis of the poverty map we identify the localities that will receive priority in the targeting scheme (we consider initially regions, then provinces, then communities, etc.). We calculate the amount a that will be targeted to all persons in the priority regions (budget S divided by the total number of poor people in the country, N_p). We simulate the targeting scheme in turn for each replication r by allocating a to all persons in our priority regions (irrespective of whether, in replication r , those regions are particularly poor or not) until the budget is exhausted. We re-calculate the post-transfer

national poverty rate in replication r . The average post-transfer national poverty rate across all replications provides our estimate of how poverty will have changed as a result of the transfer scheme. This expected poverty rate can be compared to the original estimate of national poverty from the poverty map, and to the estimate of the poverty associated with an untargeted lump-sum transfer.

Calculating “Equivalent Gains”

In thinking about the “performance” of the transfer scheme we are interested not only in the poverty impact of a specific scheme, but also in how much more “expensive” a given poverty reduction is without targeting as opposed to with geographic targeting. To explore the latter we apply a variant on the simulation procedures described above whereby we calculate how much smaller the overall budget S can be in order to achieve the same poverty impact with optimal targeting as with the untargeted uniform lump-sum transfer.

Distance from Perfect Targeting

As has been emphasized in section II, the poverty map cannot provide reliable estimates of poverty below some level of aggregation. This is because the idiosyncratic component of the overall standard error on the poverty estimate becomes progressively larger as poverty is being measured over progressively smaller groups. For policy makers therefore, the poverty map cannot be viewed as a tool to assist with the identification of, say, poor households or small groups of households. A rough rule of thumb is that poverty map estimates become unreliable for communities of less than 1000-5000 households.

It remains of interest, however, to ask how much of a further reduction in poverty could be expected if, rather than being limited to geographic targeting of communities of 1000-5000 households, policymakers could actually target individual households.¹⁵ The “dump-file” underpinning the policy simulations described above can be drawn on to shed light on this question. While this file cannot provide a reliable listing of poor households, it is possible to simulate the change in poverty in a given locality associated with household-level targeting on the basis of exactly the same optimal transfer scheme procedure described above (but where each household h now corresponds to a separate community c). In the discussion of results below we provide an assessment of the distance (in terms of poverty reduction) between *feasible* geographic targeting (given the poverty map) and the ideal of perfect household-level targeting. To the extent that this distance is large we can provide a sense of the potential benefit of combining geographic targeting with additional complementary targeting efforts.

Budgets and Poverty Lines

Before turning to a discussion of results, we describe briefly our selection of budgets, S , and our poverty lines. As has been mentioned in Section I, we assume that the budget available for distribution has been exogenously set. As is intuitively clear, the potential benefits from targeting will vary with the overall size of budget, and for this reason we conduct the simulation analyses described above for two different budget sizes. In each of the three countries examined here we identify the per capita consumption value of the 25th percentile of

¹⁵ Relatedly, one could ask how much of a further benefit could one expect if, rather than being compelled to provide only lump-sum transfers to poor communities, policymakers were able to combine geographic targeting with, say, means testing within poor communities.

the consumption distribution.¹⁶ We scale this consumption value by the total population. Our benchmark budget is set to equal 5% of this total value, and we experiment with a higher budget of 10% of this figure as well.

Gains from targeting also vary with the choice of poverty line. In this study, we select as benchmark poverty line, that line which yields a 20% headcount rate in each of our three countries.¹⁷ For comparison we experiment as well with a higher poverty line corresponding to a headcount rate of 40%.

IV. Results

Optimal targeting

Tables 2a-2c present the basic results from our simulations with the optimal targeting scheme in Ecuador, Madagascar and Cambodia. There are five main observations. First, in all three countries, the availability of disaggregated data on poverty can help to improve on a uniform lump-sum transfer across the entire population. Targeting transfers to poor localities, in accordance with the optimization scheme outlined above, yields lower values of the national FGT2 than when the budget is transferred as a uniform lump-sum transfer to the entire population. Second, the more disaggregated the poverty map, the greater the improvement over the uniform lump-sum transfer. Traditional household surveys are generally able, at best, to provide estimates of poverty at the first administrative level. The simulations here suggest that with estimates of poverty at the 2nd or 3rd administrative level, further improvements in terms of impact on the FGT2 with a given budget are attainable, and are non-negligible. Third,

¹⁶ The consumption distribution is constructed on the basis of the average, across r replications, of household-level predicted per-capita consumption in the population census.

¹⁷ Within each replication r , the predicted per-capita consumption level associated with a 20% headcount rate is identified. The average across the r replications of this predicted consumption level is then taken as poverty line. It is clear that this poverty line will

the gains from spatial disaggregation are attenuated the higher the budget and poverty line. In all three countries examined, the benefits from geographic targeting are most pronounced in our base case with the lowest budget and lowest poverty line. Fourth, while the general patterns we observe are similar across our three case-study countries, they are not identical. In particular, in Ecuador, where our focus is only on rural areas, the gains in general from targeting at the local level are somewhat more muted than in Madagascar and Cambodia. Fifth, even though our base-case low budget represents a considerable resource envelope (as evidenced by the sizable impact on poverty of even a uniform lump-sum transfer) it is clear that optimal targeting at the lowest possible level of disaggregation is far from sufficient to eliminate poverty altogether (see further below).

Table 3 considers, for rural Ecuador, the statistical precision of comparisons of poverty based on our alternative targeting scenarios. While the point estimates on the FGT2 measure generally take very low values, differences in estimated poverty across scenarios remain strongly significant (Table 3). To ascertain the statistical significance of poverty differences we return to the optimal transfer simulations and estimate not FGT2 values, but rather the difference in the estimated FGT2 based on optimal targeting at the parroquia level vis-à-vis targeting at the uniform, region, province, and cantonal level. As we can see in Table 3 not only does targeting at the parroquia level yield the lowest level of the FGT2 across our four budget/poverty line scenarios and our different levels of geographic disaggregation, but this estimate is also lower in a statistical sense. We can see that, in general, the estimated difference in point estimates is roughly 7-10 times the value of the standard error of that difference, even when we compare canton-level targeting against parroquia-level targeting.

not necessarily yield a 20% headcount rate within each replication, nor would it yield such a rate for average per capita consumption at the household level (averaging across r replications).

Tables 4a-4c present the findings in Tables 2a-2c from a different perspective. We ask now how much more cheaply one could have achieved the same impact on poverty as with the uniform lump-sum transfer, if use had been made of a poverty map in an optimal fashion. We can see, for example, in Table 4a that in rural Ecuador even a poverty map at the region level (the level of spatial disaggregation that is representative in the household survey) would have permitted the same reduction in the FGT2 as the uniform transfer (in the low budget, low poverty line base case) at only 83% of the cost of the uniform transfer. With a more detailed poverty map that allows for disaggregation down to the parroquia level, the same impact could have been achieved at only 58% of the cost of the uniform transfer. In Madagascar, and Cambodia the savings are even more striking. For example, in these two countries one would need, respectively, only 37% and 31% of the uniform transfer budget to achieve the same reduction in the FGT2 with optimal targeting at the firaiana and commune-level (Tables 4b and 4c).

“Naïve” targeting

The optimization scheme implemented above is intuitively straightforward. But working out exactly how much to give to communities is not always easy to describe. Given that the design and implementation of targeting schemes is often part of a political process, and that there is generally a need to be able to explain allocations in a simple and clear manner, it is of some interest to ask whether gains from spatially disaggregated geographic targeting are also significant when the poverty map is combined with simplistic, non-optimal, transfer schemes. We describe in Tables 5a-5c the results in our three countries of transferring our given budget on the basis of one such naive and simplistic scheme. It is striking that in all

three countries, the reduction in the FGT2 achievable with a naïve scheme, in combination with the most disaggregated poverty map, is fairly sizeable. Broadly, the reduction in the FGT2 on the basis of this scheme is similar to the impact with the optimal scheme at one level of aggregation higher. For example, in Ecuador, combining the parroquia-level poverty map with our naïve scheme (in the base case of a low budget and low poverty line) yields an estimated national level FGT2 of 0.0178, marginally higher than the 0.0177 attainable with a canton-level poverty map and an optimal targeting scheme. Similarly, in Madagascar, firaiana-level targeting with the naïve scheme yields a national FGT2 estimate of 0.0135 which can be compared to the FGT2 of 0.0138 achievable with the fivondrona-level poverty map and optimal targeting. Again, in Cambodia, commune level targeting with the naïve scheme yields an estimated FGT2 of 0.010 which is equal to that achievable with optimal targeting at the district level.

The performance of the naïve scheme becomes much less appealing when combined with a higher budget and higher poverty line. Indeed, there are cases where the impact on the FGT2 of targeting with this naïve scheme is less pronounced than with a non-targeted uniform lump-sum transfer (see for example column 2 in Tables 5a-5c and Figures 1 and 2).

Perfect targeting

How well does optimal geographic targeting at the lowest level of spatial disaggregation permitted by our poverty maps compare to the hypothetical case where we have information on the poverty status of every individual household in the country? We can answer this question in a straightforward manner by noting that the cost of eliminating poverty under the assumption of perfect targeting (i.e. it is possible to observe the precise welfare level

of every household and to tailor the transfer received by each household perfectly) is provided by the FGT1 measure of poverty (weighted by the poverty line and the total population). Thus, we can calculate from our poverty mapping database the hypothetical cost of eliminating poverty if it were possible to target the poor perfectly (and there were no behavioral responses). We can then take this budget and target it, instead, geographically, at the lowest level of geographic disaggregation that we feel that the poverty map can support. How far are we from having eliminated poverty when our transfer occurs at this geographic level rather than having been tailored to the precise circumstances of each poor household? In rural Ecuador, on the basis of the lower poverty line, optimal parroquia-level geographic targeting of this budget reduces the FGT2 from 0.028 to 0.0177, only a 37 percentage point decline. At the higher poverty line, optimal geographic targeting of the budget that, in principle, could eliminate poverty altogether, reduces the FGT2 from 0.070 to 0.032, a 54 percentage point decline. Very similar results obtain in Madagascar and Cambodia (for example, see Table 6 for Cambodia).

Why does optimal geographic targeting on the basis of a detailed poverty map fall so far short of the ideal? In a companion paper, Elbers et al (2004) analyze evidence on the level and variation of inequality within poor communities. They show that in three countries, including two of the countries examined here (Ecuador and Madagascar), within-community inequality varies widely across communities. Some communities exhibit levels of inequality as high, or higher, than at the national level, while others are significantly more equal. An important conclusion from this study is that there should be no presumption of lower levels of inequality in poor communities. In fact, in the three countries studied by Elbers et al (2004), median inequality is highest amongst the bottom quintile of communities (ranked in terms of

average per capita consumption, or in terms of the headcount rate of poverty) and this quintile also displays the highest degree of variation of inequality across communities. The implication of this finding is that within poor communities, even small ones with populations of 5000 households or less, there are likely to be both poor and non-poor households. Community level targeting that transfers a uniform amount to all individuals within these small communities is thus likely to continue to suffer from leakage. The poverty impact of such targeting will thus fall short of what would have been possible if perfect targeting were available.

V. Discussion

The recent literature on micro-estimation of welfare based on combined survey and census data offers a promising avenue for analyzing the potential poverty impact from a variety of policy proposals, such as geographically targeted transfer schemes. In this paper we have taken recently completed “poverty maps” for three countries, Ecuador, Madagascar and Cambodia, and have explored the extent to which the availability of local estimates of poverty can help to strengthen targeting performance. We have taken the raw census-level output files from the poverty mapping methodology in these three countries and have simulated the impact on poverty of transferring an exogenously given budget to geographically defined sub-groups of the population according to their relative poverty status. We have asked to what extent effectiveness of targeting in reducing poverty improves as we define these sub-groups at progressively lower levels of spatial disaggregation.

We have found large gains from targeting smaller administrative units, such as districts or villages. We have shown that the largest gains from geographic targeting occur when the overall budget available for transfer is relatively modest, and when the poverty line is not so

high as to classify most of the population as poor. We have shown further that the benefits from targeting are most clearly discerned when expressed in terms of budgetary savings of achieving a given rate of poverty reduction. We have noted, however, that despite the availability of reasonably precise estimates of poverty at the level of communities with populations of perhaps only 5000 households, transfer schemes targeting such communities are still unable to achieve the kind of success that would be attainable if household-level income or consumption data were available. We suggest therefore that a potentially useful way forward is to combine fine geographic targeting on the basis of a poverty map with within-community targeting based on either self-selection or alternative targeting mechanisms.

Our assessment of targeting performance has been based on an optimal use of estimates from poverty maps. There might be grounds for concern that the design of transfer schemes based on such optimized routines suffers from lack of transparency and would be difficult to describe in simple terms. We have considered, therefore, an alternative transfer scheme, based on a naïve, non-optimal use of the poverty map. We have found that while this naïve scheme does not achieve the same success in our three countries as the optimal targeting scheme, its performance remains surprisingly good in the case where the overall budget and poverty line are both relatively low. On the other hand, when the budget and/or poverty line is relatively high, the naïve scheme we implemented can perform less effectively than even a uniform lump-sum transfer. To the extent that policy makers are concerned about issues of communication and transparency, our results suggest that there are conditions under which even simplistic targeting schemes based on the poverty map can yield encouraging results.

We would be remiss not to restate the important caveats that attach to these results. First, we have assumed in this paper that the budget available for distribution is exogenously

determined. We abstract away entirely from the question of how the transfers are to be financed. Yet the design of targeting schemes may influence the mobilization of resources. Second, we have not addressed the possibility of variations in the costs of administering transfer schemes with the degree of disaggregation. Third, we have not taken into account the possibility that behavioral responses might vary with the degree of disaggregation in the transfer scheme. Finally, we have not examined optimal targeting in the context of local political-economy considerations or in the presence of community specific public goods. The results in this paper should thus be viewed as suggestive only, and should be combined with detailed assessments of these mitigating factors.

References

- Baker, J., and Grosh, M. (1994). Poverty Reduction Through Geographic Targeting: How Well Does It Work? *World Development* 22(7): 983-995.
- Coady, D. and Morley, S. (2003). *From Social Assistance to Social Development: Targeted Education Subsidies in Developing Countries* (Center for Global Development and International Food Policy Research Institute).
- Dasgupta, I., and Kanbur, R. (2003). Community and Anti-Poverty Targeting. Unpublished manuscript.
- De Janvry, A., Sadoulet, E., Calva, L.F., and Solaga, I. (2003). Concept Paper and Terms of Reference for the External Evaluation of SEDESOL's Microregions Program. Unpublished manuscript, Dept. of Agriculture and Resource Economics, UC Berkeley.
- Demombynes, G., Elbers, C., Lanjouw, J.O., Lanjouw, P., Mistiaen, J. and Özler, B. (2002) 'Producing an Improved Geographic Profile of Poverty: Methodology and Evidence from Three Developing Countries' in van der Hoeven, R. and Shorrocks, A. (eds) Growth, Inequality and Poverty: Prospects for Pro-Poor Economic Development (Oxford: Oxford University Press).
- Duclos, J., Makdissi, P., and Wodon, Q. (2003). Poverty-Efficient Transfer Programs: the Role of Targeting and allocation Rules. CIRPÉE Working Paper 03-05.
- Elbers, C., Lanjouw, J.O. and Lanjouw, P. (2002). Micro-Level Estimation of Welfare. Policy Research Working Paper 2911, Development Research Group, the World Bank, Washington D.C.
- Elbers, C., Lanjouw, J.O. and Lanjouw, P. (2003). Micro-Level Estimation of Poverty and Inequality. *Econometrica* 71(1): 355-64.
- Elbers, C., Lanjouw, P., Mistiaen, J., Özler, B., and Simler, K. (2004). Unequal Inequality of Poor Communities. *World Bank Economic Review* (forthcoming).
- Fujii, T. (2003). Commune-level Estimation of Poverty Measures and its Application in Cambodia. Unpublished manuscript, Dept. of Agriculture and Resource Economics, UC Berkeley.
- Gelbach, J. and Pritchett, L. (2002). Is More for the Poor Less for the Poor? The Politics of Means-Tested Targeting. *Topics in Economic Analysis and Policy*, Vol 2(1).
- Glewwe, P. (1992). Targeting Assistance to the Poor. *Journal of Development Economics* 38: 297-321.
- Hentschel, J. and Lanjouw, P. (1996). Constructing an Indicator of Consumption for the Analysis of Poverty: Principles and Illustrations with Reference to Ecuador. *LSMS Working Paper* No.124, DECRG-World Bank: Washington DC.

Hentschel, J., Lanjouw, J.O., Lanjouw, P., and Poggi, J. (2000). Combining Census and Survey Data to Trace the Spatial Dimensions of Poverty: A Case Study of Ecuador. *World Bank Economic Review* 14(1)147-65.

Kanbur, R. (1987). Measurement and Alleviation of Poverty. *IMF Staff Papers*, 34:60-85.

Mistiaen, J., Özler, B., Razafimanantena, T., and Razafindravonona, J. (2002). Putting Welfare on the Map in Madagascar. Africa Region Working Paper Series 34.

Ravallion, M. and Chao, K. (1988). Targeted Policies for Poverty Alleviation Under Imperfect Information: Algorithms and Applications. *Journal of Policy Modeling* 11: 213-24.

Ravallion, M. (1993). Poverty Alleviation Through Regional Targeting: A Case Study of Indonesia, in *The Economics of Rural Organization: Theory, Practice and Policy*, editors Braverman, A., Hoff, K. and Stiglitz, J. World Bank and Oxford University Press.

Table 1. Data Summary

	<u>Ecuador</u>	<u>Madagascar</u>	<u>Cambodia</u>
Household Survey			
Year	1994	1993-4	1997
Source	Encuesta de Condiciones de Vida (ECV)	Enquête Permanente Auprès des Ménages (EPM)	Cambodia Socio-Economic Survey (CSES)
Sample Size	4,500 Households	4,508 Households	6010
References	Hentschel and Lanjouw (1996); and Hentschel, Lanjouw, Lanjouw and Poggi (2000)	Mistiaen, Özler, Razafimanantena and Razafindravonona (2002)	Fujii (2003)
Population Census			
Year	1990	1993	1998
Coverage	About 10 million individuals in 2 million households	about 11.9 million individuals in 2.4 million households	About 11.0 million individuals in 2.15 million households

Table 2a: Rural Ecuador
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
Optimal Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.028	0.028	0.070	0.070
Uniform transfer	0.020	0.014	0.058	0.047
Region-level targeting (3 regions)	0.0188	0.0128	0.0565	0.0455
Province-level targeting (21 provinces)	0.0184	0.0125	0.0557	0.0447
Canton-level targeting (195 cantons)	0.0177	0.0119	0.0544	0.0433
Parroquia-level targeting (915 parroquias)	0.0167	0.0110	0.0528	0.0412

Table 2b: Urban and Rural Madagascar
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
Optimal Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.027	0.027	0.068	0.068
Uniform transfer	0.019	0.013	0.054	0.043
Province-level targeting (6 provinces)	0.0154	0.0097	0.0500	0.0382
Fivandrona-level targeting (111 fivandronas)	0.0138	0.0080	0.0473	0.0346
Firaisana-level targeting (1248 firaisanas)	0.0126	0.0071	0.0455	0.0327

Table 2c: Urban and Rural Cambodia
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
Optimal Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.019	0.019	0.052	0.053
Uniform transfer	0.014	0.009	0.042	0.033
Province-level targeting (44 urban plus rural provinces)	0.011	0.007	0.039	0.028
District-level targeting (180 districts)	0.010	0.006	0.036	0.026
Commune-level targeting (1594 communes)	0.009	0.005	0.034	0.023

Table 3: Rural Ecuador
Testing for Statistical Significance of Differences in Poverty Estimates
Comparing Parroquia-Targeted Estimates Against Other Levels of Targeting

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Parroquia-level targeting FGT2 estimate	0.0167	0.0110	0.0528	0.0412
Uniform – parroquia (s.e.)	0.0034 (0.00055)	0.0030 (0.00056)	0.0050 (0.0054)	0.0056 (0.00068)
Region – parroquia (s.e.)	0.0021 (0.00028)	0.0019 (0.00025)	0.0037 (0.00034)	0.0043 (0.00041)
Province - parroquia (s.e.)	0.0016 (0.00021)	0.0016 (0.00020)	0.0028 (0.00024)	0.0035 (0.00033)
Canton – parroquia (s.e.)	0.00095 (0.00012)	0.00095 (0.00013)	0.0016 (0.00014)	0.0021 (0.00021)

Table 4a: Rural Ecuador
Cost of Achieving the Uniform Transfer Impact When Using Optimal Targeting
Expressed as a Percentage of Uniform Transfer Outlay

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Uniform transfer	1.000	1.000	1.000	1.000
Region-level targeting (3 regions)	0.827	0.881	0.904	0.938
Province-level targeting (21 provinces)	0.760	0.849	0.822	0.900
Canton-level targeting (195 cantons)	0.667	0.793	0.727	0.831
Parroquia-level targeting (915 parroquias)	0.584	0.711	0.645	0.745

Table 4b: Urban and Rural Madagascar
Cost of Achieving the Uniform Transfer Impact When Using Optimal Targeting
Expressed as a Percentage of Uniform Transfer Outlay

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Uniform transfer	1.000	1.000	1.000	1.000
Province-level targeting (6 provinces)	0.607	0.699	0.704	0.777
Fivandrona-level targeting (111 fivandronas)	0.464	0.568	0.571	0.654
Firaisana-level targeting (1248 firaisanas)	0.376	0.482	0.501	0.584

Table 4c: Urban and Rural Cambodia
Cost of Achieving the Uniform Transfer Impact When Using Optimal Targeting
Expressed as a Percentage of Uniform Transfer Outlay

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Uniform transfer	1.000	1.000	1.000	1.000
Province-level targeting (44 urban plus rural provinces)	0.545	0.696	0.625	0.715
District-level targeting (180 districts)	0.414	0.579	0.402	0.616
Commune-level targeting (1594 communes)	0.308	0.419	0.436	0.513

Table 5a: Rural Ecuador
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
“Naïve” Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.028	0.028	0.070	0.070
Uniform transfer	0.020	0.014	0.058	0.047
Region-level targeting (3 regions)	0.0194	0.0145	0.0565	0.0456
Province-level targeting (21 provinces)	0.0201	0.0188	0.0559	0.0457
Canton-level targeting (195 cantons)	0.0188	0.0172	0.0552	0.0446
Parroquia-level targeting (915 parroquias)	0.0178	0.0162	0.0540	0.0425

Table 5b: Urban and Rural Madagascar
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
“Naïve” Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.027	0.027	0.068	0.068
Uniform transfer	0.019	0.013	0.054	0.043
Province-level targeting (6 provinces)	0.0172	0.0156	0.0503	0.0383
Fivandrona-level targeting (111 fivandronas)	0.0148	0.0129	0.0489	0.0363
Firaisana-level targeting (1248 firaisanas)	0.0135	0.0114	0.0480	0.0348

Table 5c: Urban and Rural Cambodia
Impact on FGT2 of Targeting at Different Levels of Geographic Disaggregation
“Naïve” Targeting Scheme

	Low Budget Low Poverty Line	High Budget Low Poverty Line	Low Budget High Poverty Line	High Budget High Poverty Line
Original FGT2	0.019	0.019	0.052	0.053
Uniform transfer	0.014	0.009	0.042	0.033
Province-level targeting (44 urban plus rural provinces)	0.012	0.011	0.039	0.030
District-level targeting (180 districts)	0.011	0.010	0.038	0.028
Commune-level targeting (1594 communes)	0.010	0.008	0.036	0.026

Figure 1: Comparing Optimal and “Naïve Targeting in Madagascar with the Low Budget and Low Poverty Line




Figure 2: Comparing Optimal and “Naïve Targeting in Madagascar with the High Budget and Low Poverty Line

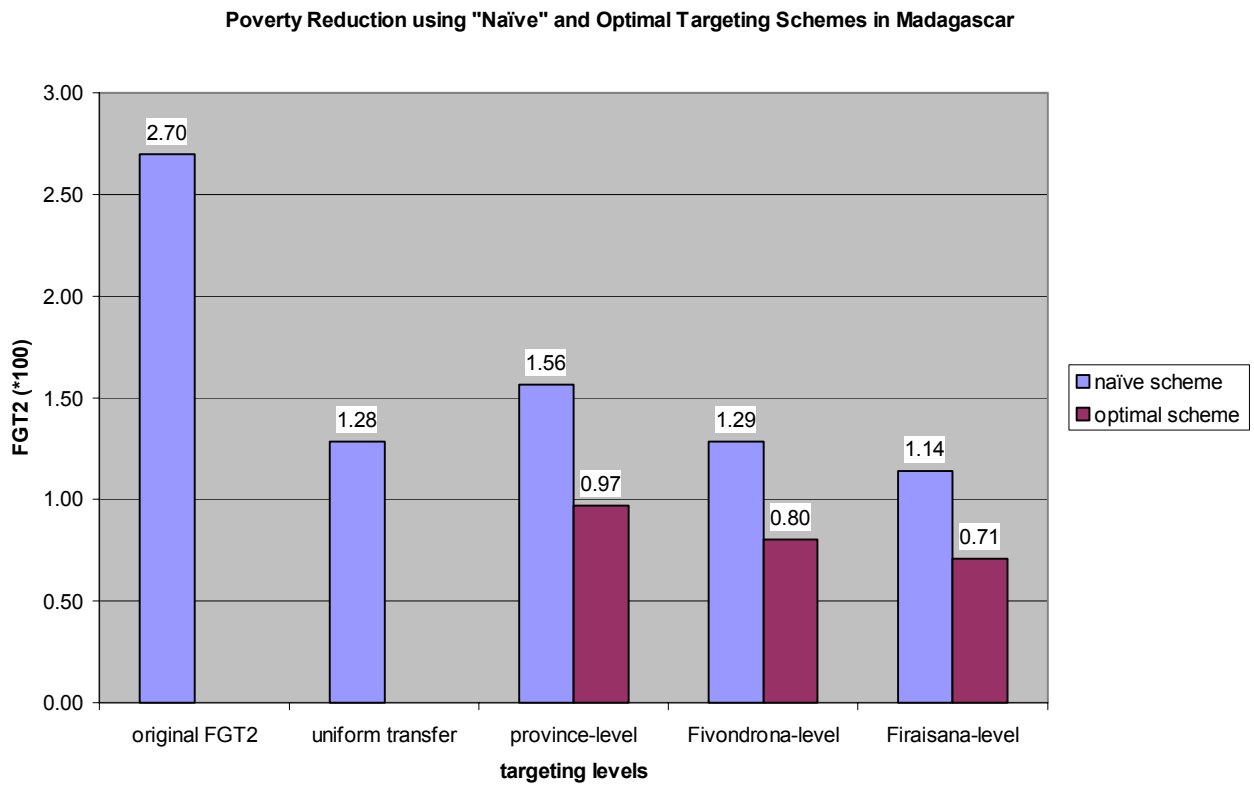


Table 6: Distance between Optimal Geographic Targeting and “Perfect” Targeting in Cambodia

	FGT2 (*100)	% spent on non-poor	% reduction in FGT2
Budget= Total Poverty Gap for <i>Low Poverty Line</i>			
<i>Level of targeting</i>			
Pre-Transfer	1.93		
Lumpsum Transfer	1.47	81.2	23.9
Province*Urban/Rural (44)	1.23	71.2	36.4
District (180)	1.12	67.0	41.9
Commune (1594)	0.99	62.5	48.7
Household (2130544)	0.00	0.00	100.0
Budget=Total Poverty Gap for <i>High Poverty Line</i>			
Pre-Transfer	5.23		
Lumpsum Transfer	2.80	64.1	46.42
Province*Urban/Rural (44)	2.31	56.0	55.86
District (180)	2.13	53.1	59.31
Commune (1594)	1.80	49.4	64.12
Household (2130544)	0.00	0.00	100.0