

CDE
May, 2006

**POVERTY TARGETING IN PUBLIC PROGRAMS: A
COMPARISON OF SOME NONPARAMETRIC TESTS
AND THEIR APPLICATION TO INDIAN
MICROFINANCE**

Isha Dewan

Email: isha@isid.ac.in

**Indian Statistical Institute
New Delhi**

Rohini Somanathan

Email: rohini@econdse.org

**Delhi School of Economics
University of Delhi**

Working Paper No. 154

Centre for Development Economics
Department of Economics, Delhi School of Economics

Poverty targeting in public programs: A comparison of some nonparametric tests and their application to Indian microfinance.*

Isha Dewan[†] and Rohini Somanathan[‡]

May 8, 2007

Abstract

Many popular social programs have limited coverage among households at the very bottom of the income and wealth distribution. If a program reaches the poor, but neglects the destitute, the (pre-program) income distribution of participants and non-participants will cross. We are interested in the statistical methods that can be used to test for this particular pattern of program participation. Our numerical simulations suggest that recently developed tests for distribution crossing are powerful even when the two distributions under study are fairly similar and they can be usefully combined with more standard quantile tests to characterize program participation among the very poor. We apply this approach to data on household expenditures and membership of micro-credit groups in India and find that participation among the poorest households in the study area was lower than that of slightly richer households.

*Thanks to the PRADAN staff for many useful interactions and their organizational support during the period of data collection, to Charlie Brown, Julie Cullen, Esther Duflo, Roger Gordon, Subhash Kochar, B.L.S. Prakasa Rao and E. Somanathan for comments, and to the Consultative Group to Assist the Poorest at the World Bank for funding the survey.

[†]Indian Statistical Institute, New Delhi, isha@isid.ac.in

[‡]Delhi School of Economics, rohini@econdse.org

1 Introduction

All public programs face the challenge of reaching intended beneficiaries. Documented deficiencies in many older social transfer mechanisms have led to the emergence of innovative methods of reducing poverty. Microfinance schemes have become especially popular and are at the centre of redistributive policies in many countries.

These programs have undoubtedly improved the lives of millions of poor households. The Grameen Bank of Bangladesh reports loans to over 7 million poor borrowers in 2007 and many similar microcredit institutions exist all over the world. Yet, there are reasons to suspect that these and other government initiatives to alleviate poverty do not adequately serve the really destitute. Extreme poverty is often accompanied by levels of education, nutrition and information that are not conducive to program participation. Social ostracism may also make it hard for some of these households to be involved in group activities and their limited contact with bureaucrats may exclude them from state subsidies.

There is some evidence to suggest the empirical validity of some of these mechanisms. Morduch (1998) finds that although microfinance institutions in Bangladesh stipulate that only households below a given asset threshold are eligible for receiving loans, these eligibility rules are often violated. In a different context, Paxson and Schady (2002) estimate benefits from social funds in Peru and find that the poorest 7% of households are less likely to benefit than those that are moderately poor.

The purpose of this paper is to explore statistical methods that are appropriate to test hypotheses about the relative exclusion of households in the bottom tail of an income or asset distribution. If a program is well-designed to cater to the poor but not to the destitute, the population distributions of participants would cross that of non-participants from below. Standard tests of first order stochastic dominance would not be applicable in such cases, and, depending on the position of the crossing, some higher-order dominance relations may also not hold.

We perform numerical simulations with alternative pairs of distributions and we find that nonparametric tests that explicitly test for distribution crossing are powerful, even when the two distributions under consideration are fairly similar. We show how these might be usefully combined with standard sign tests for income quantiles to characterize the population that is best served by a particular program.

Specifically, we use the method outlined in Chen et al. (2002) to test the null hypothesis of equal distributions against the alternative of a single crossing, with the distribution of participants crossing that of non-participants from below. This approach also provides us with an estimate of the crossing point. We use sign tests to test hypotheses about differences in the quantiles of the two distributions to the left of this estimated crossing point. Typically, the further an income quantile is from the crossing point, the greater the distance between the two distributions and the more powerful a quantile test is likely to be. The estimate of the crossing point is therefore useful in choosing an appropriate quantile at which the two distributions can be compared.

In Section 3 we present power comparisons for examples drawn from families of uniform and two-parameter exponential distributions. We find the test for distribution crossing to be extremely powerful even for small sample sizes and for cases in which the chosen distributions for the two groups are very similar. Not surprisingly, the power of the sign test increases as we test for quantiles further away from the crossing point. The simulation results lead us to conclude that if there is reason to believe, *a priori*, that there is a unique threshold below which program participation is difficult, a combination of tests for distribution crossing and quantile tests can be useful in examining the pattern of household participation in a social program.

In Section 4, we apply these methods to test for poverty targeting in a rapidly growing microfinance program in India. Our data do not include a comprehensive measure of income, but do contain annual household expenditures under several major heads such as food and clothing, health, education and entertainment. The data come from

a household survey that was administered to a random sample of members of newly formed microcredit groups (to ensure that program benefits had not affected household income) and a randomly selected sample of non-members. We compare the total household expenditure in the above categories for members and non-members. We find evidence that the distribution of members cross that of non-members from below. Sign tests indicate a preponderance of non-members among the poorest quarter of the population; in this range, expenditure quantiles for the member distribution appear to be higher than those for the population as a whole.

2 Statistical Methods

We denote the population distribution of program participants by $F(x)$ and non-participants by $G(x)$. In our microfinance application, these are the members and non-members of women's savings and credit groups. Sample sizes for these two groups are n and m respectively and the two samples are denoted by X_1, \dots, X_n and Y_1, \dots, Y_m .

We are interested in testing a null hypothesis of equal distributions against the alternative that the distribution of participants cross that of non-participants from below. The Kolmogorov Smirnov test is commonly used to test for the equality of two distributions. This has two principal disadvantages in our setting: the typically low power of this test makes it unlikely that we reject the null when the population distributions of the two groups are fairly close. Also, as is well recognized, the rejection of the null does not provide much information about the alternative.

We focus on methods which test the null of equal distributions against the explicit alternative of a single crossing. Deshpande and Shanubhogue (1989), Hawkins and Kochar (1991) and Chen et al. (2002) present alternative tests of this type. We choose the test statistic proposed by Chen et al. (2002) over the Deshpande and Shanubhogue

test because unlike the latter test, it requires no prior information on the position of the crossing point and such prior knowledge is unlikely in most applications. Their test is also shown to be more powerful than the Hawkins and Kochar test in a variety of examples and has the additional advantage of providing an estimate of the crossing point. Since by definition, the fractions of both groups are equal at this point, this is an estimated upper bound on the incomes of the households that are relatively neglected by the program. For the sake of completeness, we briefly summarize their methodology before proceeding to discuss our results.

Let Z_1, \dots, Z_N be the combined sample of X 's and Y 's and $Z_{(1)} < Z_{(2)} < \dots, Z_{(N)}$ be the order statistics of this sample. We wish to test

$$H_0 : F(x) = G(x)$$

against the alternative

$$H_A : F(x) < G(x) \text{ when } x < x^* \text{ and } G(x) < F(x) \text{ when } x > x^*.$$

for some unknown income level x^* .

Consider the function

$$\lambda(x) = \sup_{t \leq x} (G(t) - F(t)) + \sup_{x \leq t} (F(t) - G(t)) - |F(x) - G(x)|$$

The value of $\lambda(x)$ at any point x is simply the sum of the largest difference between the distribution of non-participants and participants for values less than x , the largest difference between participants and nonparticipants above x , and the distance between these functions at x . Under the null hypothesis, $\lambda(x)$ is zero and under the alternative, it is maximized at the crossing point.

The sample counterpart of $\lambda(x)$ is given by

$$\lambda_N(x) = \sup_{t \leq x} (G_m(t) - F_n(t)) + \sup_{x \leq t} (F_n(t) - G_m(t)) - |F_n(x) - G_m(x)|$$

where $F_n(x)$ and $G_m(x)$ are empirical distribution functions corresponding to $F(x)$ and $G(x)$ respectively. The test statistic is based on the largest sample value of $\lambda_N(x)$:

$$\sup_x \lambda_N(x) = \max_{0 \leq j \leq N} \lambda_N(Z_{(j)})$$

The order statistic $Z_{(j)}$ is an estimate of the income level at which the distributions cross. Chen et. al (2002) propose the statistic

$$J_N = \sqrt{\frac{mn}{N}} \max_{0 \leq j \leq N} \lambda_N(Z_{(j)}) \quad (1)$$

for testing H_0 against H_A . Large values of the statistic are unlikely under the null of equal distributions. Exact critical points for small samples are tabulated in Chen et al. (1998). The asymptotic distribution is not standard and the authors present asymptotic critical regions based on Monte-Carlo simulations. The relevant critical value is presented with our simulation results below.

In the following sections we combine the above approach with sign tests for population quantiles to the left of the estimated crossing point. Sign tests are probably the most commonly used nonparametric procedure to test for population quantiles.¹ If two distributions cross in the manner specified under the alternative hypothesis stated above, income quantiles for the participant distribution would be higher than those for non-participant distribution to the left of the crossing point. Moreover, we might expect to find the difference between the two quantiles to be largest as we move away from the crossing point to smaller income levels. For different values of p , we test the null hypothesis that the p^{th} quantiles of both distributions are equal by using a test statistic based on the number of participants below the $p(m+n)^{th}$ ordered observation of the combined sample. We use asymptotic critical values based on the normal distribution.

¹A description of these tests can be found in most standard statistical texts. See for example, Gibbons and Chakraborti, 1992.

3 Simulation Results

We compare the test procedures described above using simulated data from three alternative pairs of distributions for participants and non-participants. In each case, we start with a base distribution of non-participants $G(x)$ and alternative distributions of participants $F(x)$, all of which cross $G(x)$ at the median income level. We apply the crossing point test to samples from the base distribution and each variant of the participant distribution, and then use sign tests to test hypotheses about differences in quantiles of the two distributions to the left of the estimated crossing point.

We consider three different families of distributions. In the first case, $G(x)$ is uniformly distributed on $(0, 1)$ and $F(x)$ has uniform distribution on (a, b) . The point of intersection of the two distributions is $\frac{a}{1+a-b}$. In order to ensure that the two distributions cross at the median level of income, we choose our parameters a and b such that $a + b = 1$. We choose alternative values of a and b and examine the power of the test when the participant distribution comes closer to the base distribution. Our




Figure 1: $G(x) \sim U(0, 1)$, $F(x) \sim U(a, b)$.

next two cases are drawn from the family of two-parameter exponential distributions.

$$\begin{aligned} G(x) &= 1 - \exp[-\lambda_2(x - \theta_2)], \quad x \geq \theta_2, \lambda_2 > 0 \\ F(x) &= 1 - \exp[-\lambda_1(x - \theta_1)], \quad x \geq \theta_1, \lambda_1 > 0 \end{aligned}$$

The two distributions cross at $x^* = \frac{\lambda_1\theta_1 - \lambda_2\theta_2}{\lambda_1 - \lambda_2}$ which is median income if

$$\frac{\lambda_1\lambda_2(\theta_1 - \theta_2)}{\lambda_1 - \lambda_2} = \ln 2. \quad (2)$$

We first set the scale parameter $\lambda_1 = 1$ and the location parameter $\theta_1 = 0$ for the non-participant distribution $G(x)$, and, for the distribution $F(x)$, we consider various departures in location and scale parameters subject to (2). In the second case, we set $\lambda_1 = 2$ and $\theta_1 = 0$ for $G(x)$. These two cases are shown in Figures 2 and 3 respectively. We choose parameters in Case 2 for which the two distributions are very close. We do this in order to examine the behavior of these tests for when the null and alternative hypotheses are very similar.

For each of the cases described above, we use equal-sized samples of 2 sizes (i) $n = m = 50$ and (ii) $n = m = 100$ and perform 5,000 iterations. For the crossing test we use the simulated 5% critical point of 1.529 given in Chen et al.(2002). We perform sign tests for quantiles of order $p = .1, .2, .3$ and $.4$ and use 5% asymptotic critical values based on the standard normal distribution. Power comparisons of these tests based on these simulations are shown in Tables 1- 3.

The tables show that the power of all the tests increases as the sample size increases. For all the cases we consider, the power of the distribution crossing test is uniformly higher than that of the sign test. The sign test is powerful for quantiles of small order but declines quite rapidly as we consider quantiles close to the crossing point (the median). Not surprisingly, the power of all tests increases as we move away from the base distribution, and the difference between the null and the alternative hypotheses becomes more marked.




Figure 2: $G(x) \sim \exp(1,0)$




Figure 3: $G \sim \exp(2,0)$

Table 1: Power Comparisons for Uniform Distributions.

Value of a	Sample size per group	Test Power				
		Sign Test for p^{th} quantile				Crossing Test
		p=.1	p=.2	p=.3	p=.4	
.4	50	1	.989	.770	.343	1
	100	1	1	.934	.405	1
.3	50	.992	.864	.425	.173	1
	100	.999	.988	.698	.466	1
.2	50	.902	.494	.156	.081	.981
	100	.995	.701	.280	.065	1
.1	50	.347	.141	.055	.036	.466
	100	.575	.174	.068	.021	.836

Table 2: Power comparisons for Exponential Distributions, Case 1.

Value of λ	Sample size per group	Test Power				
		Sign Test for p^{th} quantile				Crossing Test
		p=.1	p=.2	p=.3	p=.4	
3	50	1	.966	.660	.247	1
	100	1	.998	.864	.272	1
2.5	50	1	.914	.493	.189	1
	100	1	.998	.788	.201	1
2	50	.996	.827	.350	.122	.988
	100	1	.957	.593	.146	1
1.5	50	.896	.490	.149	.074	.809
	100	.989	.700	.273	.057	.989

Table 3: Power comparisons for Exponential Distributions, Case 2.

Value of λ	Sample size per group	Test Power				
		Sign Test for p^{th} quantile				Crossing Test
		p=.1	p=.2	p=.3	p=.4	
4	50	.994	.842	.349	.120	.995
	100	1	.971	.610	.123	1
3.5	50	.972	.713	.258	.097	.964
	100	1	.897	.443	.093	1
3	50	.888	.490	.162	.060	.792
	100	.994	.695	.26	.055	.985
2.5	50	.537	.190	.062	.047	.394
	100	.797	.284	.097	.031	.723

4 An Application to Data from an Indian Micro-finance Program.

4.1 Data

Our data is from a microfinance program in the state of Jharkhand in Central India. The growth of microfinance in India has been quite different from that in most other countries in that it has been dominated by non-government organizations, voluntary savings groups and nationalized commercial banks rather than by specialized micro-finance institutions (Harper, 2002). Non-government organizations organize women into *self-help groups* with between 10 and 20 members. These groups initially pool their own weekly savings and facilitate risk-sharing among their members by lending accumulated savings to those members in need of credit. After doing this successfully for about a year, the group opens a savings account at a nearby commercial bank

and its members collectively borrow from the bank for both income-generating activities and for their consumption needs. The Reserve Bank of India (India's central bank), has issued guidelines to nationalized commercial banks that encourage bank linkages with these groups and the banks are offered credit at subsidized rates to further promote such lending.

The microfinance program we study is administered by PRADAN, a non-government organization which has created about 2,000 groups in the state of Jharkhand since 1992. Jharkhand is among the poorest of the 27 Indian states, with over half its population below the national poverty line. Our main objective was to examine the extent of participation in the program by households at the very bottom of the income distribution. Our strategy was to compare the distribution of household spending across samples of randomly selected households who had members in newly formed *self help groups* and other households in the area.

Our sampling frame consisted of households in 100 villages in which at least one new group had been formed during the period April 1st to June 30th, 2002. A total of 149 groups had been formed in these villages during the specified period. The 100 villages with new microcredit groups were partitioned into 4 geographical clusters and a simple random sample of 6 villages was drawn from each cluster. A total of 24 respondents were surveyed from each of these villages- 6 of them members of microcredit groups in the village and the remaining 18 randomly selected non-members from the same village. Our sample therefore consists of 576 households in 24 villages.

The survey was conducted over a period of two months starting in August 2002. Very little lending takes place in the months immediately following group formation and a comparison of the characteristics of households in the program with those of other randomly chosen households in the area can therefore be used to evaluate the extent to which the program targeted the poor.

We collected data on a large number of economic indicators such as the quantity

and type of food consumed, annual expenditures under several major heads, the size and condition of the household's main dwelling, land owned and cultivated and the possession of durable goods. We have no direct data on household incomes and instead use the sum of annual expenditures for several major expenditure categories as a proxy for income. These categories are: Clothing and footwear, schooling, health, renovation, entertainment and other social expenditures and food and non-food expenditures at the weekly village market (converted to annual figures). We also asked respondents about their contact with the government bureaucracy and about any benefits received from government sponsored programs. Responses to these questions allow us to examine the extent to which poor households who do not participate in the program receive assistance from other official poverty-alleviation programs.

If we compare households with new group members to those with no members, we find mean values of most variables are very similar for the two groups. The member households own a little more land on average, but have slightly lower household expenditures. Based on standard tests for differences in means we reject the null hypothesis of equal means for only one of the listed variables, the number of meals consumed in the two days prior to the survey. Of the households with members in self-help groups 17% had fewer than three meals a day for the two days prior to the survey, whereas this is true for 33% of the non-member households. This type of food inadequacy is a particularly stark characteristic of extreme poverty and these numbers suggest that the program might neglect or be ill-suited to households in the bottom tail of the income distribution. The slightly higher mean expenditures for non-member households and the higher fraction of households with food inadequacy are consistent with participation rates being lower at the two extremes of the income distribution than in the middle, i.e. with the hypothesis of the member distribution crossing the non-member distribution from below.

The empirical distributions of annual household expenditures for the two groups are shown in Figure 4. Based on the methods described in Section 2, we test the null hypothesis that the distributions are equal against the alternative that the distribution

of members crosses that of non-members from below. The test statistic, J_N takes a value of 2.04 and we therefore reject the null hypothesis at the 1% level. The level of annual household expenditure at the estimated crossing point is 12,240 Indian rupees, or 290 in U.S. dollars at the current exchange rate.

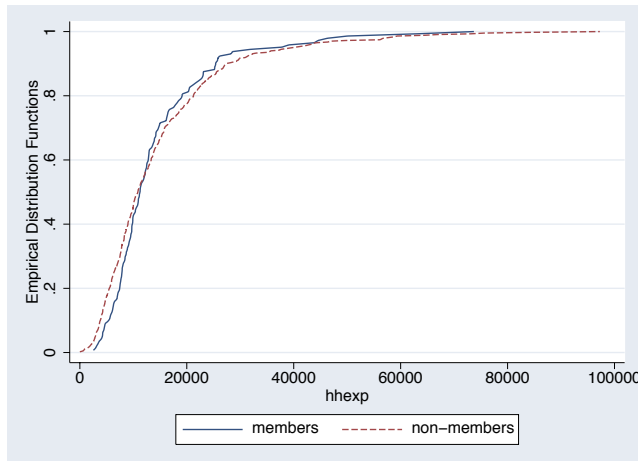


Figure 4: Empirical distribution functions of annual household expenditure for members and non-members.

We find that 56% of each group in our sample has expenditures below this point. If the probability of being included in the program rises with income and then falls, the crossing point gives us an upper bound on the income levels of the households neglected by the program, since we have equal fractions of members and non-members below this income-level. To get a better idea of exactly who is neglected, we could supplement this by other methods, such as kernel density estimates or proportions at different population quantiles below this level. Figure 5 presents kernel densities of household expenditures for the two groups.²

The estimated densities first intersect at annual expenditure level of 6,432 rupees. This is about \$145 or about \$24 per capita given the average household size of 6. A

²We use Epanechnikov kernels and optimal bandwidths. Half bandwidths equal 2415 for non-member households and 2153 for member households.

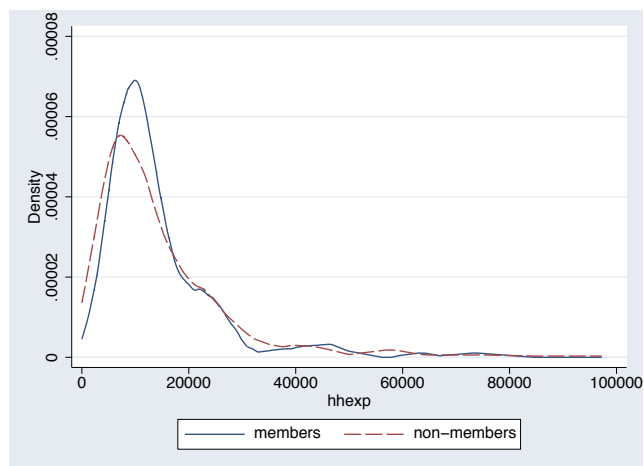


Figure 5: Kernel densities of annual household expenditure for members and non-members.

little under 23% of households are below this level of expenditure. Table 4 compares household characteristics for the first expenditure quartile with the rest of the sample. The first-quartile households consume about 40% fewer foodgrains (by weight) as the other households, they spend about two-thirds less on clothing and footwear, they live in smaller dwellings and eat fewer meals. The households with low participation in the microcredit program also seem to be excluded from other public programs and the political process more generally. Only 11% of these households had ever approached a government official compared to 32% of other households. Perhaps the most striking observation is that the fraction of households receiving subsidized foodgrains from a government anti-poverty program was no different for these households than for the rest. It appears that poor households with low participation rates in microfinance programs also have lower access to other welfare programs.

Table 4: Household characteristics of the first quartile and other households.

	first expenditure quartile	Top 3 expenditure quartiles
Meals consumed during the two days prior to the survey	5.16	5.48
Number of rooms in dwelling	2.23	3.59
Foodgrain consumption per day in normal times (kilograms per household)	2.46	4.37
Annual expenditure on clothing and footwear (rupees)	1216	4029
Land owned (hectares)	.83	1.22
Fraction ever approached government official	.11	.32
Fraction received goverment subsidized foodgrains	.51	.51

5 Conclusions

For a variety of plausible reasons, poverty alleviation programs may not reach the poorest households in an area. We explore how statistical methods can be used to test the hypothesis that a program neglects the bottom tail of a distribution in favor of slightly richer households.

We find that recently developed tests of distribution crossing are extremely powerful in testing such a hypothesis even when differences in the distributions of the two groups (participants and non-participants) are not large. We advocate that these be combined with quantile tests to examine program participation among very poor households.

We apply these methods to a microfinance program in India. We find evidence that the population distributions of program participants and non-participants cross, and households in the bottom income quartile are less likely to be program participants than those in the middle two quartiles. These households also appear to have limited access to public programs which are, in principle, designed for their benefit: they are no more likely to be on official poverty lists and government-sponsored social programs.

References

- [1] Chen, G.J., J.H. Chen and Y.M. Chen, *Statistical inference on comparing two distribution functions with a possible crossing point*, Technical Report, Department of Mathematics, Anhui University, 1998.
- [2] Chen, Guijing, Jiahua Chen and Yuming Chen, “Statistical inference on comparing two distribution functions with a possible crossing point”, *Statistics and Probability Letters*, 60, 2002, 329-341.

- [3] Deshpande, J.V. , Shanubhogue, A. “Test for two sample nonparametric dispersion alternatives.”, Technical report, University of California, Santa Barbara, 1989.
- [4] Gibbons, Jean Dickinson and Subhabrata Chakraborti, *Nonparametric Statistical Inference*, M. Dekker, New York, 1992.
- [5] Hawkins, D.L and Subhash Kochar, “Inference for the Crossing Point of Two Continuous CDF’s”, *The Annals of Statistics*, 19(3), 1991, 1626-1638.
- [6] Harper, Malcolm, “Grameen Bank Groups and Self-help Groups; what are the differences?”, <http://www.alternative-finance.org.uk>, 2002.
- [7] Morduch, Jonathan, *Does Microfinance Really Help the Poor? New Evidence from Flagship Programs in Bangladesh*, working paper, 1998.
- [8] Paxson, Christina and Norbert R. Schady, “The Allocation and Impact of Social Funds: Spending on School Infrastructure in Peru”, *World Bank Economic Review*, 16(2), 2002, 297-319.
- [9] Somanathan, Rohini, “Poverty Targeting in PRADAN’s Microfinance Program: A Study Conducted on Behalf of CGAP ”, available from the Consultative Group to assist the Poorest at the World Bank, Washington D.C., April, 2003.