Benefit Incidence, Public Spending Reforms, and the Timing of Program Capture

Peter Lanjouw and Martin Ravallion

Assessments of the distributional effects of public spending reforms have generally been based on average rates of program participation by income or expenditure group. This practice can be deceptive because the socioeconomic composition of participants can change as a social program expands or contracts. The geographic variation found in 1993–94 household survey data for rural India is used to estimate the marginal odds of participating in schooling and antipoverty programs. The results suggest early capture of these programs by the nonpoor. It is shown that conventional methods for assessing benefit incidence underestimate the gains to the poor from higher public outlays and underestimate the losses from cuts.

Benefit incidence analysis is widely used to assess the distributional impact of public spending. Typically, the average participation rate for each public program is tabulated against household income or expenditure per capita, using data from a household survey. The public subsidy rate for each program is then applied to the participation rates to determine the incidence of program spending—to assess, for example, whether the poor gain more than the rich. Discussions of the method and examples for developing countries can be found in Meerman (1979), Selowsky (1979), Meesook (1984), Hammer, Nabi, and Cercone (1995), Selden and Wasylenko (1995), van de Walle (1995, 1998), and Demery (1997).

This article examines whether this now-standard methodology provides a reliable guide to the distributional impact of public spending reforms. These reforms typically entail marginal changes in spending across one or more programs. But benefit incidence calculations are based on averages. To see why this difference matters, consider a publicly supplied private good, such as subsidized schooling or food. Unlike a private good obtained in a competitive market, one cannot buy or sell as much of a publicly supplied good as one wants—consumers are quantity constrained. And unlike a pure public good (in which everyone faces the same quantity constraint), the way in which a program's outlays are allocated

across consumers comes into play. That allocation is typically the outcome of a political process. The distributional impact of a change in supply will then depend on the abilities of different socioeconomic groups to influence that political process. Those abilities, in turn, will depend in part on the history of allocations made under the program at the time reforms commence.

In these circumstances the averages used in traditional benefit incidence analysis can be deceptive in predicting the marginal changes that would arise from public spending reforms. For example, suppose that the nonpoor were able to capture most of the benefits when the program was first introduced but are now virtually satiated at the margin. Then the poor will gain a large share of the extra benefits from an expansion of the program and may lose heavily from its contraction, even though they receive a small share of the average benefits.

We use household survey data on participation in primary school and the main antipoverty programs in rural India. First, we estimate average participation rates, using standard methods. We then estimate marginal participation rates, exploiting the fact that there are large differences across Indian states in the scale of each type of public program. Thus we are able to compare the average and marginal odds of participation and so test for bias in estimates of the distributional impact of spending changes based on standard benefit incidence analysis.

I. Does the Composition of Program Participation Vary with Program Size?

The population is divided into two or more groups according to consumption expenditure (or some other welfare indicator). For each expenditure group we know the participation rate in a public program. We examine the effect of changing the overall size of the program as part of a public spending reform. If the group-specific participation rates stay the same, then the composition of program participation is said to be *homogeneous* (strictly, homogeneous of degree one). This means that if, say, 40 percent of participants are poor when 1,000 participants are covered, then 40 percent are also poor when 2,000 are covered. Policy conclusions drawn from standard benefit incidence analyses implicitly assume homogeneity. But there is no obvious reason why this should hold.

Nonhomogeneous participation can arise when the poor are able to capture program benefits at certain times in the program's history, but not others. This can occur even when the program is ostensibly targeted solely to the poor. Either the government is unable to target perfectly—because of information or incentive problems—or targeting only the poor is not a political equilibrium in that the government relies on the support of the nonpoor.[1] The timing of program capture will depend on how the costs and benefits of participation vary with the scale of the program. Social programs invariably impose costs on participants. These costs could take the form of taxes or fees for financing the programs. Or they could be hidden,

1. For an overview of these issues, see Besley and Kanbur (1993); for a model of the political economy of targeting in which perfect targeting is not an equilibrium, see Gelbach and Pritchett (1997).

deadweight losses arising from financing methods or from participation itself. Such losses include the opportunity cost to parents of children's time in school or the cost to a nonpoor person of illegally securing participation in a means-tested program. These costs are likely to vary with the scale of the program.

For example, the geography of program placement can generate nonhomogeneous participation. Consider a country in which poor areas tend to be more remote and hence less convenient for program staff to reach. Initial placement is in less remote areas. When the program is first set up, the poor will find it more costly to participate than when the program eventually expands into remote areas. The nonpoor will be able to capture the program early because it is more accessible to them. But after some point marginal gains start to favor the poor.

General equilibrium effects could also produce rising costs of participation that differ between the poor and the nonpoor. For example, although a small public works program may not affect wages in alternative work, a large program may bid up wages and hence increase the expected forgone income of program participants. To the extent that the nonpoor face better chances of getting work, they will find the public employment program less and less attractive as it expands. Again, early capture by the nonpoor can be expected.

Late capture is also possible. For example, it may be far easier for the (theoretically ineligible) nonpoor to bribe officials to gain access to the program once it is widely available, when a nonpoor participant would be less conspicuous.

A Model

We illustrate these arguments more formally in a simple political economy model. The model assumes that the government wants to reduce poverty, but that it faces a political constraint in that it cannot impose a welfare loss on the nonpoor. The program does impose costs on the nonpoor (such as taxes for financing the program) that depend on average participation of the poor and nonpoor. Let the cost to the nonpoor be $C(X)$, where X is average participation and C is a positive, smoothly increasing function with $C(0) = 0$. (Other factors may enter this cost function, such as the proportion of the population who is not poor, but these can be ignored for the purpose of our model.) Marginal cost is $C'(X)$, which could either increase or decrease with X (depending on whether C is convex or concave). The program's benefits are allocated between nonpoor households, who participate at the rate X_n, and poor households, who participate at the rate X_p. The corresponding per capita benefits are $B(X_n)$ and $B(X_p)$, where the function B is increasing from $B(0) = 0$. (To simplify the notation, we assume that B is the same for the poor and the nonpoor. Allowing the functions to differ does not affect our analysis, although it could affect the assessment of benefit incidence in practice.) The utility of a nonpoor household is $U[Y_n + B(X_n) - C(X)]$, where Y_n is the household's exogenous income, and the function U is strictly increasing.

Political feasibility requires that the nonpoor do not lose from the program. In other words, a necessary condition for the program to continue is that:

(1) $U[Y_n + B(X_n) - C(X)] \geq U(Y_n),$

where $U(Y_n)$ is the utility of the nonpoor without the program $(X_n = X = 0)$. Since the government values gains to the poor, the political economy constraint in equation 1 will be binding in equilibrium. (If it were not binding, there would be a politically feasible change that benefited the poor.) Because equation 1 must then hold with equality, we have:

(2) $B(X_n) = C(X).$

Solving equation 2 for

(3) $X_n = \Psi(X)$

tells us how program participation by the nonpoor varies with average participation, given the political economy constraint. The participation rate of the poor is also a function of X, namely:

(4) $X_p = [X - N_n \Psi(X)] / N_p,$

where N_n and N_p are the proportions of the population who are nonpoor and poor, respectively. The marginal change in participation by the nonpoor as the program expands—the marginal participation rate of the nonpoor—is given by

(5) $\Psi'(X) = C'(X) / B'(X_n).$

To see what this model implies for the timing of program capture, consider first the special case in which B is linear (constant marginal benefits). Then, it is plain from equation 5 that X_n will be concave (convex) in X—implying that the marginal gains to the poor tend to rise (fall) as the program expands—whenever the marginal cost to the nonpoor is decreasing (increasing). Early (late) capture will occur when the cost function is concave (convex).

With declining marginal benefits to the nonpoor (B is concave), a convex cost function still implies late capture. For early capture the cost function must be sufficiently concave. Differentiating equation 5 with respect to X, we can see that

(6) $\Psi''(X) = [C''(X) B'(X_n) - X_n'(X) C'(X) B''(X_n)] / B'(X_n)^2.$

For early capture, $\Psi''(X) < 0$, to be the only politically feasible option in this model, it is necessary and sufficient that the (absolute) elasticity of marginal cost, $-XC''(X) / C'(X)$, exceed the elasticity of participation by the nonpoor, $X\Psi'(X) / X_n$, times the elasticity of the marginal benefit from the program's allocation, $-X_n B''(X_n) / B'(X_n)$.

Early versus Late Capture

The above model illustrates how, for public programs with relatively large start-up costs, early capture by the nonpoor may be the only politically feasible option, particularly when start-up costs must be financed domestically. For example, in exchange for paying taxes to cover these costs, the nonpoor may de-

mand a sizable share of the initial benefits, such as by requiring that the program not be located in inaccessible, poor areas. Only later, when the marginal costs of program expansion are lower, will it be politically feasible to reach the poor.

Figure 1 illustrates the case of early capture. The figure shows the group-specific participation rate as it varies with the average rate, that is, the function $\Psi(X)$ and $2X - \Psi(X)$ for the nonpoor and poor, respectively. (For convenience, the figure is drawn assuming that there are equal numbers of poor and nonpoor.) The nonpoor capture the bulk of the gains initially but become progressively satiated. Imagine we are at point A, where the poor and nonpoor are participating equally. Given the average participation rates, the standard benefit incidence analysis would conclude that expanding the program would not benefit the poor relative to the nonpoor. This conclusion is plainly wrong: most of the gains from a small aggregate expansion at point A (to, say, point B) would go to the poor. Similarly, a small cut in the program at point A would be borne mostly by the poor.

Now consider the case of late capture as illustrated in figure 2. Suppose we are initially at the average participation rate A. Although the poor are participating more than the nonpoor, it is the nonpoor who would capture most of the gains from increasing the level of average participation (to, say, point B) and incur most of the loss from retrenchment.

We would expect some public programs to be more like the early capture model and others to be more like the late capture model. For example, it is likely that children of better-off parents will be the first to gain from public spending
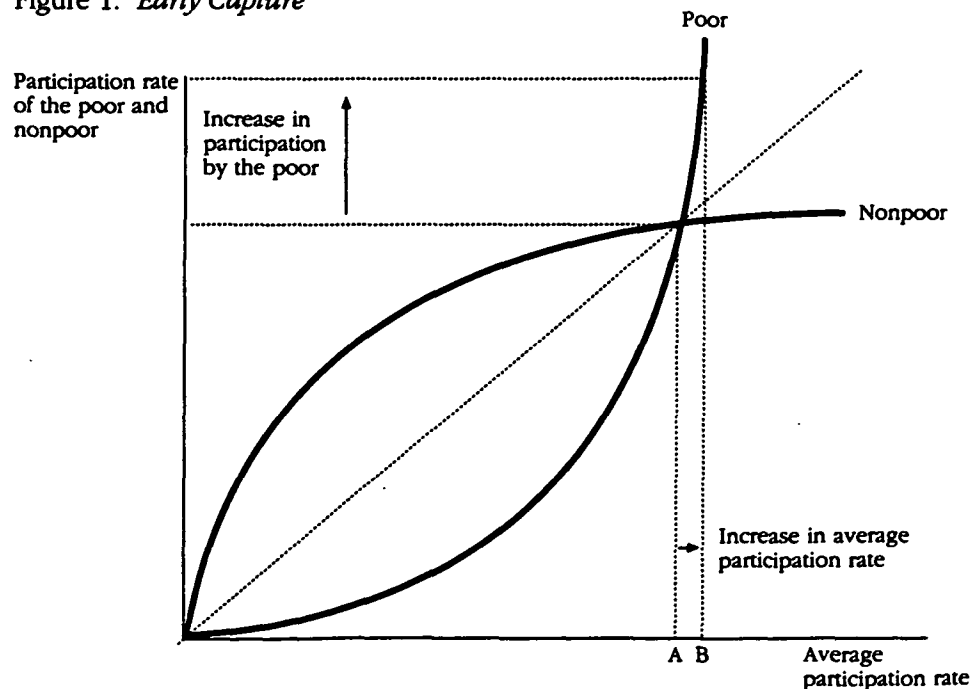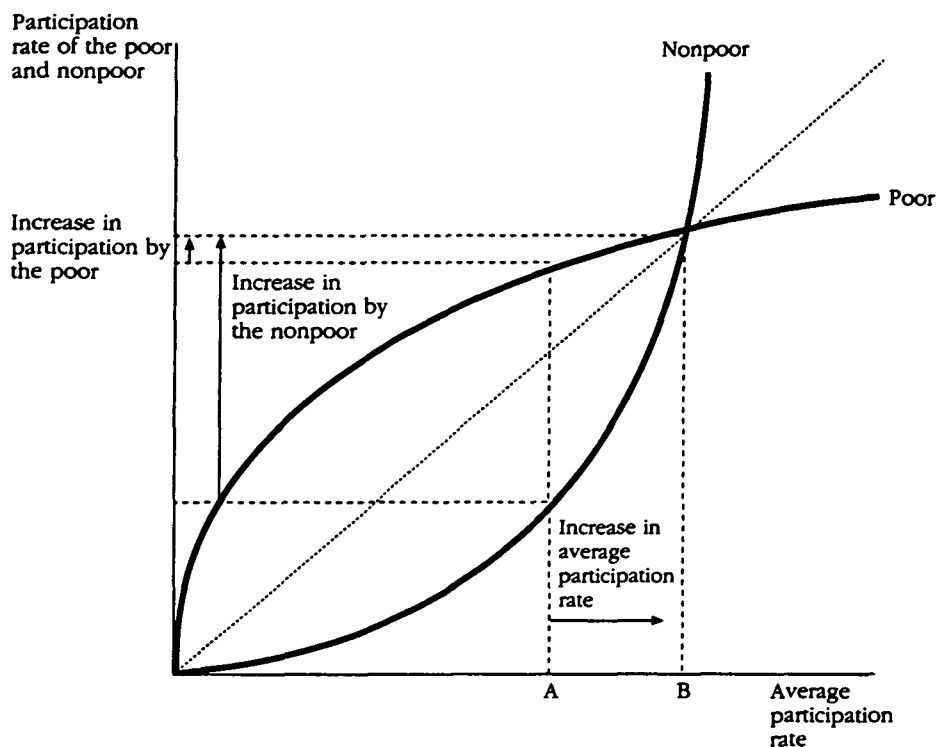
Figure 1. *Early Capture*

Figure 2. *Late Capture*

on education (early capture). But they will become satiated in due course, at which point marginal gains will go to the poor. By contrast, consider a food rationing scheme that is initially targeted to the poor. In time, political pressures to favor middle-income groups may lead to higher marginal gains for the nonpoor (late capture).

This discussion suggests that the political economy of program capture may contain important clues to some poorly understood issues concerning the welfare impacts of changes in public spending. One such issue is whether or not there are politically feasible ways of protecting the poor from cuts in social spending. Ravallion (forthcoming) explores this issue further, providing evidence that spending cuts were borne more heavily by the poor in an antipoverty program in Argentina.

The political economy of program capture can also help us to understand the empirical relationship between intercountry differences in public spending on social programs and aggregate outcome indicators. Aggregate data often show this relationship to be weak. But Bidani and Ravallion (1997) have developed an econometric specification that allows them to compare the effects of differences in countries' public health spending on health indicators for the poor and the nonpoor. They find that differences in health spending matter far more to the

health outcomes of the poor than of the nonpoor. This is what we would expect if the nonpoor capture the inframarginal gains.

II. Measuring Participation Rates

The average participation rate is the proportion of households in a given expenditure or income quintile that participates in the program. The average odds ratio of participation (herein, the average odds of participation) is defined as the ratio of the participation rate of one quintile to the overall average. The marginal odds ratio of participation (herein, the marginal odds of participation) is the increment in the program participation rate of a given quintile when there is a change in aggregate participation. Differences between the marginal and average odds of participation reflect differences in the incidence of inframarginal spending. Only if participation is homogeneous will the two be everywhere the same.

The average odds of participation can be calculated from the survey data in a straightforward way. How can we estimate the marginal odds of participation? We have only a single cross-sectional survey (as is typically used in benefit incidence analysis), including data on program participation across geographic areas ("regions") within states. We can readily calculate the average participation rate for a given program for each quintile and each region. The participation rate for a given quintile varies across regions according to the level of public spending on the program in the state to which each region belongs. To estimate the marginal odds of participation by program and expenditure quintile, we can regress the quintile-specific participation rates across regions on the state's average participation rate (for all quintiles and all regions) for each program.[2]

Ordinary least squares will give a biased estimate of the marginal odds of participation, because the specific region and quintile participation rates (on the left side) are implicitly included when calculating the state's overall mean participation rate (on the right side). To deal with this problem, we use the "leave-out mean" as an instrumental variable for the state's average participation rate. The leave-out mean is the mean for the state excluding the specific region and quintile participation rates that correspond to each observation in the data. For example, if we are using the data for quintile three in region five within state ten, then the leave-out mean is the average for all regions and quintiles within state ten, excluding quintile three in region five.

How can we interpret the marginal odds of participation? As with average participation rates, to infer overall incidence we must also know the program's subsidy rate. In conventional benefit incidence analysis the subsidy rate for each program is typically assumed to be constant across geographic areas and income groups. For example, it is assumed that the cost to the government is the same if

---

2. Our method appears to be new, although models of outcome indicators stratified by socioeconomic group are familiar from past work. Deolalikar (1995), for example, studies the cross-sectional differences in health outcomes for poor and nonpoor children in Indonesia.

a poor person participates or if a rich person participates. Using that assumption, we can infer from the marginal odds of participation how an increase in public spending on a given program will affect each quintile. We will be able to make partial tests of that assumption.

III. DATA

Our analysis is based on the household-level data from India's National Sample Survey (NSS) for 1993–94. This survey includes standard data on consumption expenditures, demographics, and educational attainment, including school enrollment. In addition, this round of the NSS also asked about participation in three antipoverty programs: public works schemes, a means-tested credit scheme called the Integrated Rural Development Programme (IRDP), and a food rationing scheme called the Public Distribution System (PDS). We collate data on participation in these programs with data on total consumption expenditures per person at the household level. A household is said to have participated in a public works program if any household member worked for at least 60 days on public works in the preceding 365 days. A household participated in the IRDP program if it received any assistance from IRDP in the past five years. And a household participated in the PDS program if it purchased any commodity from a ration or fair price shop in the past 30 days.

We ranked sampled households by total consumption expenditure per person (including imputed values of consumption from own production) normalized by state-specific poverty lines. Quintiles are defined over the entire rural population, with an equal number of people in each. Thus the poorest quintile includes the poorest 20 percent of the nation's rural population in terms of consumption expenditure per capita.

These data are not ideal. The relationship between participating in the IRDP over the past five years and consumption expenditure over the past month may be a poor indication of the program's incidence because participants' living standards may have changed considerably over five years. There are also concerns about the adequacy of participation as an indicator of use of the PDS. For example, the rich may buy only a small quantity of the rationed good (although this conjecture is not consistent with other data on the incidence of PDS purchases; see Radhakrishna and Subbarao 1997). Another possibility is that the individual participant may have a different standard of living than the household as a whole. In the case of public works projects the data will likely include people who participate in public works projects but are not part of antipoverty programs.

The sample size (for rural areas) of the 1993–94 NSS is 61,464 households. We conduct the analysis at the level of the NSS region, of which there are 62 in India, spanning 19 states. Each NSS region belongs to only one state. So, in the basic model, for any given combination of quintile and program, we regress the sample participation rates from the 62 regions on the average participation rate (irrespective of quintile) from each of the 19 states.

Recall that in benefit incidence analysis predicting the incidence of changes in public spending from the estimated marginal odds of participation requires the assumption that the average subsidy rate (given participation) is constant. For public works programs and the IRDP there is no obvious way in which the subsidy rate conditional on participation could vary by household expenditure per person within a given state. But we can expect variation between states. For the PDS income effects on demand for the rationed goods could create differences in the subsidy rate across quintiles within a given state.

We are able to test the assumption of a constant subsidy rate for public works programs, the IRDP, and primary school enrollment, but we do not have the data to do so for the PDS. For each program we regress the state's per capita spending on its average participation rate plus four of the quintile-specific participation rates. We are unable to reject the null hypothesis that the parameter estimates on the quintile-specific participation rates are jointly zero. The probability values for the F-tests are, respectively, 0.57, 0.11, and 0.23 for primary schooling, public works, and the IRDP. The coefficients on a state's average participation rate are highly significant, as we would expect. Thus we find no evidence that the subsidy rate varies significantly by quintile. This helps us to justify the constant-subsidy assumption when interpreting our results.

IV. RESULTS

We begin with primary school enrollment for children ages 5–9 years (table 1). Our calculated enrollment rates from the NSS are appreciably lower than those obtained from schools themselves, on which official enrollment rates are based. The official primary enrollment rate for India was higher than 100 percent in 1993. Although there are differences in definition—for example, we have confined attention to the age group 5–9 and so excluded late starters—there are reasons to believe that biases in official sources lead to overestimation of enrollments in India (Kingdon 1996).

Enrollment rates rise with household expenditures per capita nationally and in all states, and they tend to be higher for boys than for girls. (Lanjouw and Ravallion 1998 give full results by state.) But there are marked differences among states. In Kerala, for example, there is less difference among the quintiles and between boys and girls (indeed, enrollment rates are slightly higher for girls from the poorest quintile) than in Bihar or Punjab.

The average odds of enrollment suggest that subsidies to primary schools mildly favor the nonpoor. Notice, however, that we cannot split public and private schooling in the data; public school enrollment may be lower for the nonpoor. While the average enrollment rate is higher for the richest quintile, the relationship between region-specific enrollment rates and states' average rates is steeper for the poorest quintile (Lanjouw and Ravallion 1998 give scatter plots). Thus the marginal odds of participation are higher for the poor, even though their average participation rate is lower.

Table 1. *Average Primary School Enrollment in Rural India*

	Boys		Girls		Total	
Quintile	Enrollment rate (percent)	Average odds of enrollment (mean = 1.0)	Enrollment rate (percent)	Average odds of enrollment (mean = 1.0)	Enrollment rate (percent)	Average odds of enrollment (mean = 1.0)
1 (poorest)	42.6	0.75	31.6	0.66	37.2	0.71
2	53.4	0.93	43.1	0.91	48.6	0.90
3	60.5	1.07	50.3	1.06	55.8	1.08
4	66.1	1.16	58.6	1.26	62.6	1.21
5	69.9	1.23	65.2	1.38	67.7	1.31

Note: The table gives the average primary school enrollment rates as a percentage of children aged 5–9 and the average odds of enrollment, defined as the ratio of the quintile-specific enrollment rate to the mean enrollment rate. Households are ranked by total expenditure per person in forming the quintiles.

Source: Authors' calculations based on India's 1993–94 National Sample Survey.

We estimate the marginal odds of being enrolled by regressing the participation rates of each quintile across regions on the states' average participation rates (table 2). The numbers in table 2 can be interpreted as the gain in subsidy incidence per capita for each quintile from a one-rupee increase in aggregate spending on each program. For example, if an extra 100 rupees per capita is spent on primary schools, public expenditures per capita going to the poorest quintile will rise by 110 rupees. These are instrumental variables estimates in which the leave-out mean is the instrument for the state average participation rate.

The estimates of the marginal odds of participation suggest that expanding primary schooling would be decidedly propoor at the margin. (As in standard benefit incidence analysis, future earnings gains from better education are not factored into this calculation.) The implication for the incidence of subsidies to primary education is clear (given our inability to reject the constant unit-subsidy assumption). The average odds of participation given in table 1 suggest that the share of the total subsidy going to the poorest quintile is only 14 percent (0.71 times one-fifth). By contrast, the marginal odds of participation, given in table 2, imply that the poorest quintile would obtain about 22 percent of an increase in the total subsidy going to primary education.

There is also a gender difference between the average and marginal odds of participation. The average odds of poor children attending school are higher for boys (0.75 compared with 0.66 for girls). However, the marginal odds are almost identical (1.09 compared with 1.08). These results are clearly not consistent with homogeneous participation. Marginal gains from expanding primary schooling in rural India are much better distributed than average gains.

Turning now to the antipoverty programs, we see that for both public works programs and the IRDP participation rates fall as expenditures per person rise (table 3). But the rate of decline is not large; the odds of the poorest quintile participating in public works programs are 1.23 compared with 0.83 for the

Table 2. *Marginal Odds of Primary School Enrollment in Rural India*

Quintile	Boys	Girls	Total
1 (poorest)	1.09	1.08	1.10
	(6.90)	(9.65)	(8.99)
2	0.91	0.91	0.97
	(6.05)	(6.99)	(7.92)
3	0.92	0.84	0.87
	(5.85)	(6.54)	(7.65)
4	0.66	0.66	0.67
	(4.10)	(4.28)	(4.77)
5	0.53	0.70	0.67
	(4.08)	(5.53)	(5.69)

Note: The table gives the instrumental variables estimates of the regression coefficients of the quintile-specific primary school enrollment rates across regions on the average rate by state. The leave-out mean enrollment rate is the instrument for the actual mean. The numbers in parentheses are *t*-ratios.

Source: Authors' calculations based on the 1993–94 National Sample Survey.

Table 3. *Average Participation Rates for India's Main Antipoverty Programs in Rural Areas*

Quintile	Public works programs		Integrated Rural Development Program		Public Distribution System	
	Participation rate (percent)	*Average odds of participation (mean = 1.0)*	*Participation rate (percent)*	*Average odds of participation (mean = 1.0)*	*Participation rate (percent)*	*Average odds of participation (mean = 1.0)*
1 (poorest)	5.0	1.23	6.5	1.03	69.5	0.92
2	4.6	1.13	7.1	1.13	76.7	1.01
3	4.2	1.04	6.4	1.03	77.9	1.03
4	3.5	0.86	6.0	0.96	78.1	1.00
5	3.4	0.83	5.6	0.89	76.1	1.00

Note: The table gives the average participation rates and the average odds of participation, defined as the ratio of the quintile-specific participation rate to the mean participation rate for each program.

Source: Authors' calculations based on the 1993–94 National Sample Survey.

richest quintile. The rate of decline is even lower for the IRDP. Participation rates among the richest 20 percent of the population are high even for public works programs. For the PDS the participation rate is lowest for the poorest quintile and highest for the second-richest quintile.

Keep in mind that these figures are national aggregates. We find large differences among states (full details are available in Lanjouw and Ravallion 1998). In Orissa, for example, the proportion of households in the poorest quintile participating in public works programs is more than four times that of the richest quintile; the odds of the poorest quintile participating are 1.6, well above the national mean, 1.23 (table 3). In Maharashtra the odds of the poorest quintile participating in public works programs are also well above the national average. At the other extreme, in states such as Andhra Pradesh, Gujarat, Kerala, and Tamil Nadu, the poorest quintile has lower than average participation rates.

The marginal odds of participation for the poorest quintile are highest for public works programs, while the IRDP dominates for the three middle quintiles; the marginal odds of participation for the richest quintile are higher for the PDS (table 4). (The regional plots for all programs and the poorest and richest quintiles are available from the authors.) The estimated marginal odds of participation broadly confirm the conclusion drawn from the average odds of participation—public works programs are best at reaching the poorest, while the IRDP is more effective at reaching the middle quintiles, including those living at India's poverty line (at roughly the fortieth percentile).

The difference between the marginal odds of participation for any two programs gives the estimated gain from transferring one rupee between them. For example, transferring 100 rupees per capita from the PDS to public works programs would raise public spending per capita on the poorest quintile by 10 rupees (116 − 106 = 10, using the basic model).

Table 4. *Marginal Odds of Participation for India's Main Antipoverty Programs in Rural Areas*

Quintile	Public works programs	Integrated Rural Development Program	Public Distribution System
1 (poorest)	1.16	1.11	1.06
	(3.27)	(15.49)	(8.14)
2	0.93	1.28	0.99
	(3.64)	(17.73)	(7.26)
3	0.80	1.21	0.91
	(2.98)	(23.52)	(6.88)
4	0.92	0.96	0.86
	(4.32)	(19.09)	(7.16)
5	0.55	0.39	0.81
	(3.29)	(8.06)	(6.27)

Note: The table gives the instrumental variables estimates of the regression coefficients of the quintile-specific program participation rates across regions on the average rate by state for that program. The leave-out mean participation rate is the instrument for the actual mean. The numbers in parentheses are *t*-ratios.
Source: Authors' calculations based on the 1993–94 National Sample Survey.

For both the public works programs and the IRDP it is notable that the marginal odds of participation tend to fall more rapidly moving from the poorest to the richest quintile than do the average odds. Thus the average odds of participation underestimate how propoor an increase in average spending on each of these programs will be. This difference is particularly strong for the IRDP: the average odds of participation are only slightly higher for the poorest quintile than for the richest (1.03 and 0.89), whereas the marginal odds are much higher for the poorest quintile than for the richest (1.11 and 0.39). Compared with the average odds of participation, the marginal odds of participation raise the share of total IRDP spending imputed to the poorest 40 percent of the population by 11 percent, while that imputed to the richest 20 percent falls by 56 percent. For the PDS, however, there is less difference between the average and marginal odds, so the former are a better guide to PDS incidence relative to the other programs.

As with primary school enrollment, these results are inconsistent with the homogeneity assumption. Unlike schooling, for the antipoverty programs studied here, both average and marginal odds of participation tend to be higher for the poor. But, like schooling, marginal gains from these programs tend to be better distributed than average gains.

V. CAVEATS

It is worth reviewing some of the assumptions that underpin our efforts to estimate the marginal incidence of spending on these programs. We estimate the marginal odds of participation by regressing quintile-specific participation rates across regions on the state's average participation rate (all quintiles, all regions) for each program. We do not include any other explanatory variables (such as state-level poverty rates). To the extent that other variables affect quintile-specific participation rates via their influence on states' average participation rates, they are not of concern, because it is the effect of expansion in the overall size of the programs that we are interested in evaluating. There is, however, one way in which our specification may be unsatisfactory. In the first section we outlined how political economy factors could influence program incidence by determining the timing of program capture. Yet by not including political economy variables as separate explanatory variables in our regression, we implicitly assume that they are identical across states or vary in ways that are uncorrelated with state-level average participation rates.

We are unable to control for regional fixed effects in our estimations because we do not have time series data. But we are able to examine the extent to which interstate differences in average participation rates account for interstate differences in quintile-specific participation rates. First, for each program we reestimate each model by regressing quintile- and region-specific participation rates on a full set of state dummy variables. (The effect of the state participation rate is not identified, because it is predicted perfectly by the state dummy variables.) We then compare the R^2 values from these regressions to those that we obtained from our regressions

on state average participation rates. We find that in most cases the R^2 values from the state participation rate specifications are 70 percent or more of those from the state fixed-effect regressions. This suggests that states' average participation rates capture a large share of the variance in the dependent variable attributable to state effects. In the case of primary schooling of boys the R^2 values from our specification decline from about 75 percent of the R^2 values of the state fixed-effect specification for the lowest two quintiles to an average of about 45 percent for the top two quintiles. In the case of public works programs the ratio of R^2 values averages about 50 percent for the three lowest quintiles, rises to 70 percent for the fourth, and then declines to 31 percent for the top quintile.

Second, we examine the residuals from our regression using states' average participation rates to see whether, for any given state and quintile, the average of the residuals across regions is significantly higher or lower than that observed for other states. For example, we ask whether the participation rate of the bottom quintile in Kerala or West Bengal (both of which have had long periods of left-wing governments) is unusually high given the state's average participation rate, reflecting a difference in political economy.

We find no obvious patterns in the residuals. In very few cases (looking at the average residuals per state for each of the quintile-specific and program-specific regressions) does the state's average residual exceed in absolute value the standard error of the regression as a whole. And in the few cases in which this does occur, there is no discernible pattern showing that one state appears to be consistently more effective in reaching a particular quintile across programs. The only pattern that does emerge is for primary school enrollment in both Haryana and Punjab (for boys, girls, and the full sample): the average residuals for the bottom quintile are uniformly negative and larger than one standard error. Assuming that the political economy in these two states is appreciably different, we drop them. Our estimate of the marginal incidence of additional education spending on the poorest quintile is slightly higher than that given in table 2; the marginal odds of participation are 1.16 for boys, 1.12 for girls, and 1.13 overall. For the second poorest quintile the marginal odds of participation are 0.98 for boys, 0.99 for girls, and 1.01 overall. The direction of change strengthens our main result for the comparison of average and marginal odds of participation.

VI. CONCLUSIONS

We have used a simple model of the political economy of the timing of program capture to argue that conventional benefit incidence analysis can be deceptive about the distributional impacts of public spending returns. Motivated by this model, we used regional data for India to study how the composition of program participation varies with the size of a social program. This provided a relatively simple method of estimating the marginal odds of program participation. The method can be implemented with the same basic data used in conventional benefit incidence analysis.

Our results for India suggest that average participation rates are not a reliable guide to the distributional impacts of changes in aggregate public outlays or reallocations among programs. Our estimates of the marginal odds of participation broadly confirm the qualitative conclusion drawn from the average participation numbers for the three poverty programs that we studied. However, the average odds of participation greatly understate how propoor extra spending on either public works programs or the means-tested credit scheme is likely to be. Similarly, conventional methods underestimate the loss to the poor from program cuts. The average odds also underestimate how propoor a switch from, say, the Public Distribution System to public works programs would be in India.

In the case of primary schooling the average odds of participation give the wrong qualitative result. Although the average odds of enrollment rise with expenditures per person, the marginal odds fall sharply, indicating that aggregate expansion is decidedly propoor. Indeed, the marginal odds suggest that higher subsidies to primary education are about as propoor as the best programs directed (explicitly) at fighting poverty.

For both primary schooling and poverty programs (except the food rationing scheme) our results are more consistent with the early capture model than with the late capture model. The geographic pattern of participation suggests that the nonpoor tend to be the first to gain when a program is introduced, but that high marginal gains to the poor emerge later.

These findings are tentative. The fact that we had to rely on a single cross-sectional survey meant that we were not able to eliminate the possibility of omitted state-level effects that influence distributional outcomes and are correlated with average program participation rates. Geographic panel data on program participation would allow more robust tests.

To the extent that further work supports our findings, serious doubts are raised about assumptions routinely made in discussions of the distributional impacts of social programs. The timing of program capture can mean that the poor obtain larger gains from extra spending, and are hurt more by cuts, than data on average participation rates would suggest.

REFERENCES

The word "processed" describes informally reproduced works that may not be commonly available through library systems.

Besley, Timothy, and Ravi Kanbur. 1993. "Principles of Targeting." In Michael Lipton and Jacques van der Gaag, eds., *Including the Poor*. Washington, D.C.: World Bank.

Bidani, Benu, and Martin Ravallion. 1997. "Decomposing Social Indicators using Distributional Data." *Journal of Econometrics* 77(1):125–40.

Demery, Lionel. 1997. "Benefit Incidence Analysis." World Bank, Institutional and Social Policy, Africa Technical Family, Washington, D.C. Processed.

Deolalikar, Anil B. 1995. "Government Health Spending in Indonesia: Impacts on Children in Different Economic Groups." In Dominique van de Walle and Kimberly Nead,

eds., *Public Spending and the Poor: Theory and Evidence.* Baltimore, Md.: Johns Hopkins University Press.

Gelbach, Jonah, and Lant Pritchett. 1997. "Redistribution in a Political Economy: Leakier Can Be Better." World Bank, Development Research Group, Washington, D.C. Processed.

Hammer, Jeffrey S., Ijaz Nabi, and James A. Cercone. 1995. "Distributional Effects of Social Sector Expenditures in Malaysia, 1974–89." In Dominique van de Walle and Kimberly Nead, eds., *Public Spending and the Poor: Theory and Evidence.* Baltimore, Md.: Johns Hopkins University Press.

Kingdon, Geeta G. 1996. "Private Schooling in India: Size, Nature, and Equity Effects." Working Paper 72. Development Economics Research Programme, London School of Economics. Processed.

Lanjouw, Peter, and Martin Ravallion. 1998. "Benefit Incidence and the Timing of Program Capture." Policy Research Working Paper 1956. World Bank, Development Research Group, Washington, D.C. Processed.

Lipton, Michael, and Martin Ravallion. 1995. "Poverty and Policy." In Jere Behrman and T. N. Srinivasan, eds., *Handbook of Development Economics.* Vol. 3. Amsterdam: North-Holland.

Meerman, Jacob. 1979. *Public Expenditure in Malaysia: Who Benefits and Why.* New York: Oxford University Press.

Meesook, Oey A. 1984. *Financing and Equity in the Social Sectors in Indonesia.* World Bank Staff Working Paper 703. Washington, D.C.: World Bank.

Radhakrishna, Rokkam, and Kalanidhi Subbarao. 1997. "India's Public Distribution System: A National and International Perspective." World Bank Discussion Paper 380. World Bank, Poverty Reduction and Economic Management Network, Washington, D.C. Processed.

Ravallion, Martin. Forthcoming. "Is More Targeting Consistent with Less Spending?" *International Tax and Public Finance.*

Selden, Thomas M., and Michael J. Wasylenko. 1995. "Measuring the Distributional Effects of Public Education in Peru." In Dominique van de Walle and Kimberly Nead, eds., *Public Spending and the Poor: Theory and Evidence.* Baltimore, Md.: Johns Hopkins University Press.

Selowsky, Marcelo. 1979. *Who Benefits from Government Expenditures? A Case Study of Colombia.* New York: Oxford University Press.

van de Walle, Dominique. 1995. "The Distribution of Subsidies through Public Health Services in Indonesia, 1978–87." In Dominique van de Walle and Kimberly Nead, eds., *Public Spending and the Poor: Theory and Evidence.* Baltimore, Md.: Johns Hopkins University Press.

———. 1998. "Assessing the Welfare Impacts of Public Spending." *World Development* 26(3):365–79.