Chapter 3

Behavioral Incidence Analysis of Public Spending and Social Programs

Dominique van de Walle

3.1. Introduction

The ways in which participants and other agents respond to a program can matter greatly to its distributional outcomes. For example, recipients of a transfer payment may change their labor supply or savings choices, such that the net income gain is less than the amount of the transfer. Or there may be ways in which the behavior of intervening agents — such as local governments in decentralized programs — entails that the incidence of a change in aggregate spending differs from the average incidence.

This chapter discusses some simple tools for introducing behavioral responses into the analysis of the incidence of public spending or policy changes. The chapter looks at ways of incorporating two quite distinct types of responses: firstly, those of the direct participants and the people they interact with, and secondly, the responses of administrative or political agents.

The approaches to incorporating behavioral responses reviewed in this chapter tend to be ex post in that they study interventions that have already occurred, though often drawing implications for future policies. They also tend to be non-structural, in that they do not trace through all the behavioral interlinkages that may be involved, but focus instead on the final "reduced form" relationships between outcomes and interventions. More structural approaches to *predicting* the marginal incidence of reforms not yet implemented are discussed in chapter 6.

3.2. Assumptions about behavioral responses do matter

The key issue for all incidence analysis is how to define the counter-factual of what the income (or other welfare indicator) of beneficiaries would be in the absence of the program. Only then can one determine how individuals should be ranked so as to infer program incidence. It is only by seeing the incidence of benefits according to how poor people would have been without them that we can know their distributional impact.

So, an appropriate indicator is needed to identify the poor. In conventional benefit incidence analysis of cash transfers or in-kind transfers whose cash value has been imputed, the without- intervention position is often assumed to be given by the welfare indicator (e.g. expenditures per capita) less the monetary value of the benefits secured from the publicly-provided good or program under study. This makes the strong assumption that there is no replacement through household behavioral responses. Surely there are underlying behavioral impacts of the benefits one is assigning? By the same token, the opposite assumption — treating post-transfer consumption as the welfare indicator for assessing incidence — is equally suspect. Ideally, one would like to subtract the intervention amount but add in the replacement income households would have achieved through their behavioral responses had they not benefited from the intervention.

Note that in the case of in-kind programs for which no imputed value has been included in consumption or income aggregates (as is commonly the case for public education and health programs), one does not of course have to net them out in calculating the without-intervention position. However, the issue of behavioral responses also arises for such programs; consumption or income may well be very different in the absence of publicly provided health or education for example.

The assumptions made about behavioral responses can matter greatly to the conclusions one draws from any benefit incidence study. Naturally, this is an empirical issue. Table 3.1 highlights the potential sensitivity of the incidence of average mean per capita transfers in one case, namely Yemen in 1998. The table gives incidence results under two assumptions — fully excluding or fully including transfer incomes when assigning households to pre-intervention deciles. When deciles are defined net of transfers, the results suggest that transfers are well targeted to the very poorest households. The opposite conclusion is reached when deciles are instead defined on the basis of post-transfer expenditures: transfer income is concentrated in the richest decile. Conclusions about targeting and incidence clearly depend on how the counterfactual is defined.

Table 3.1: Distribution of net public and private transfers in Yemen in 1998 under different assumptions about the propensity to consume out of transfers *(annual YR per capita)***.**

Welfare indicator:	Per capita expenditures net of transfers (net mean per capita transfers)			Per capita expenditures with tranfers fully included (net mean per capita transfers)		
1998 National deciles	Rural	Urban	National	Rural	Urban	National
1	14757	32942	17347	1181	1651	1233
2	3169	5482	3552	1625	2055	1696
3	2290	4165	2671	1650	2468	1818
4	2158	3925	2528	2331	2311	2327
5	2237	2718	2346	1985	3200	2252
6	985	2601	1352	3246	3693	3350
7	1777	3153	2106	3039	4658	3443
8	1294	3172	1780	5138	4948	5090
9	1475	3987	2146	4860	6400	5288
10	1749	2023	1851	10777	11915	11217
Total	3358	5139	3770	3358	5139	3770

Source: van de Walle (2002a) using the 1998 Yemen Republic Household Budget Survey.
Note: Deciles are formed by ranking the population by household per capita expenditures under different assumptions about the propensity to consume out of transfers. Net transfers are calculated from income and expenditure on transfers identified in the HBS—namely, income from zakat, retirement and pensions, local and foreign remittances and payments from government organizations minus transfers given on Zakat, aid to dependents, other gifts and donations. Total household expenditures includes spending on transfers, so that only transfer income needs to be netted out to get at the "net" amounts.

3.2.1 Estimating the consumption impacts of programs

There are a number of reasons for why the current consumption gains to a participant can differ from the monetary value of a program's benefit level. The program can affect savings, labor supply, and schooling choices, and it can also affect private transfers received. Without identifying the precise "structural" channel, the most direct approach to incorporating such responses is to see how much consumption changes when benefits are received.

Recognizing the importance of the behavioral responses (as illustrated in Table 3.1), a few studies have explored the issue using panel data. An example can be found in Ravallion et al. (1995) who estimate the marginal propensity to consume out of social income (PCSI) using panel data for Hungary.[1] This is then used to determine the net gain to consumption from social

[1]The PCSI can also be estimated using regression methods on a cross-section but, depending on how well one can control for heterogeneity, the results are likely to be biased by omitted variables. A third approach is to use propensity score matching with a single difference estimator (as in Jalan and Ravallion, 2002). The advantage of the latter is that a model or structure does not need to be imposed.

transfers and to construct the counterfactual consumption level without intervention. This allows a behavioral incidence analysis.

In a similar vein, van de Walle (2002c) estimates the PCSI for Viet Nam where household surveys for 1993 and 1998 contain a panel of 4308 households. Consumption of household *i* at time *t* (*t*=1993, 1998) (C_{it}) is assumed to be represented as an additive function of public transfers (T_{it}),observed household characteristics (X_{it}), time varying (δ_t) and time invariant (η_i) latent factors:

$$C_{it} = \alpha + \beta T_{it} + \gamma X_{it} + \eta_i + \delta_t + \varepsilon_{it} \tag{1}$$

There are a number of potential problems with estimating β with this equation. An endogeneity concern arises if, as is likely, transfers are correlated with time invariant household characteristics (cov(T_{it} η_i)≠0), as could result from purposive targeting to the long term poor. Endogeneity also arises if transfers are correlated with time varying determinants of consumption (cov(T_{it} δ_t)≠0 or cov(T_{it} ε_{it})≠0). This would occur if transfers target those who suffered a shock or simply because of transfer eligibility changes, such as if a pension-receiving elderly household member dies.

There is likely to be heterogeneity in the behavioral response across households. Differences in the impact of the transfers associated with observable differences in the characteristics of individuals can be introduced by adding appropriate interaction effects in equation (1), so it takes the form:

$$C_{it} = \alpha + (\beta_0 + \beta_0 X_{it})T_{it} + \gamma X_{it} + \eta_i + \delta_t + \varepsilon_{it} \tag{2}$$

One can also readily introduce random differences in impacts not correlated with the program assignment. There are also non-parametric methods (that do not need to postulate a parametric regression equation for the outcome variable); these methods are reviewed in Chapter 5. However, for the purpose of the present exposition, attention is confined to the simplest parametric model in (1).

A double differencing model where all variables are expressed in first differences, purges the estimate of fixed effects and thus deals with the first source of endogeneity. Equation (1) becomes:

$$\Delta C_{it} = \beta \Delta T_{it} + \gamma \Delta X_{it} + \Delta \delta_t + \Delta \varepsilon_{it} \tag{3}$$

With only two rounds of data, the term $\Delta\delta_t$ becomes an ordinary intercept term in a regression of the change in consumption on the change in transfers.

In the Vietnam example, this regression was initially run assuming that $\gamma\Delta X_{it}$ = 0 (characteristics don't change or don't have any effect), giving the standard "double difference" estimate of the consumption impact of transfers. This gives a β estimate of 0.45 with a heteroscedasticity and clustering-corrected *t*-statistic of 4.3. A number of different regressions are run that control for time varying household characteristics, the possibility that there are omitted variables that alter over time and affect transfers (using an instrumental estimator), and to test for possible heterogeneity in impacts. However, none of the PCSI estimates are significantly different from the initial simple double difference estimate 0.5 (van de Walle 2002c).

The study thus uses consumption expenditures net of half of the value of transfer receipts as its ranking welfare indicator. Interestingly, though perhaps completely coincidentally, the Hungary study also estimates a marginal propensity to consume out of transfer incomes of 0.5 (Ravallion et al. 1995) and, in a slightly different context, Jalan and Ravallion (2002) estimate about 50% income replacement for public transfers in Argentina.

These examples have been for cash transfer programs. However the same points apply to in-kind transfers such as publicly-provided health or education. Then one would model consumption or income as a function of participation in such programs. The same issues of endogeneity bias naturally arise, and the panel data methods described above offer an approach to addressing these issues.

3.3. Marginal incidence analysis

Another example of a behavioral incidence analysis is what is sometimes called "marginal incidence analysis," where one measures the incidence of actual increases or proposed cuts in program spending. This approach departs from standard benefit incidence analysis that attempts to estimate how the average benefits from public spending are distributed at one point in time. The latter can be deceptive about how changes in public expenditures will be distributed. It is possible, for example, that the political economy of incidence entails that the rich tend to receive a large share of the infra-marginal subsidies, while the poor benefit most from extra spending. Ravallion (1999) provides a model of the political economy of fiscal adjustment that can generate such an outcome.

3.3.1 Using single cross-sectional data to infer marginal incidence

The simplest way to identify marginal incidence is to compare average incidence across geographic areas with different degrees of program size. This is essentially the method of Lanjouw and Ravallion (1999) who used data from India's National Sample Survey (NSS) for 1993-94. This survey includes standard data on consumption expenditures, demographics and education attainments, including school enrolments. This particular NSS round also asked about participation in three key anti-poverty programs: public works schemes, a means-tested credit scheme called the "Integrated Rural Development Programme" (IRDP), and a food rationing scheme, the "Public Distribution System" (PDS). The data on participation in these programs can be collated with data on total consumption expenditure per person at the household level.

Sampled households in the NSS were ranked by total consumption expenditure per person normalized by state-specific poverty lines. Quintiles were then defined over the entire rural population, with equal numbers of people in each. So the poorest quintile refers to the poorest 20% of the national rural population in terms of consumption per capita.[2] The analysis was done at the level of the NSS region, of which there are 62 in India, spanning 19 states and with each NSS region belonging to only one state. So, for any given combination of quintile and program, the participation rates across the 62 NSS regions were regressed on the average participation by state (irrespective of quintile). The results provide estimates of the marginal incidence of participation across quintiles and indicate that expansion of primary schooling would be very pro-poor in contrast to average incidence figures that suggest the opposite. With regard to the poverty schemes, additional spending would be significantly more pro-poor than suggested by the average incidence of participation.

Although the technique requires a cross-sectional household survey only, it must contain information on program participation at household level and sufficient regional disaggregation and variance in participation for estimation to be possible. The main concern with using a single cross-sectional survey is that there may be important state-level differences in the propensity to reach the poor that are correlated with levels of social spending. One way to get at marginal incidence more robustly to such latent heterogeneity in local political factors is to assess incidence at two or more dates. This can be done using two cross-sections or panel data (if it is available) on households or regions.

[2] Note that, although this study does not allow for behavioral responses on the part of individuals and households in determining the welfare ranking indicator as above, there is nothing that prevents it from doing so.

3.3.2. Marginal incidence analysis using repeated cross-sectional data

Take the example of spending on education, and the case where two consecutive cross-sectional surveys are available with information on households with children attending school in both years. Each enrolled child is assumed to receive the same subsidy in a schooling level i. The change in quantile specific participation in education between the two years can then be represented by

$$\frac{E_{ij2}}{E_{i2}} - \frac{E_{ij1}}{E_{i1}} \qquad j = 1, 2, \ldots \qquad (4)$$

where E_{ijt} is the number of children in a given level of schooling i in welfare quantile j, at date t = 1,2 and E_{it} is total school enrollment in that level at date t. Alternatively this can be interpreted as the marginal incidence of spending on education between the two years where enrolments are multiplied by the appropriate subsidy level as in chapter 2. In the following the same simplified notation E is used to refer to both representations. E_{jt}/E_t can be interpreted as the share of total enrolments or education spending that goes to quantile j through the school attendance of its children. The expression in (4) tells us the change in each quantile's share of enrollments or spending. Alternatively, one might want to know the share of a given quantile in the total change in enrollments or education spending as given by:

$$\frac{E_{j2} - E_{j1}}{E_2 - E_1} \qquad (5)$$

The above approach can be applied to health care, social transfers, and other public spending programs for which participation at household or individual level can be identified and — if one wants to identify public spending amounts — a benefit value attributed. The important point is that there may be a big difference between average incidence at a point of time — as indicated by E_{j1}/E_1 or E_{j2}/E_2 — and the marginal incidence defined by (4) or (5).

A number of studies have examined ex-post marginal incidence in this way. Early examples looked at whether changes during the 1980s were pro-poor for Indonesia's public health sector and Malaysia's health and education sectors (van de Walle 1994; Hammer et al. 1995, respectively) and found that they were. Another study looked at the changing incidence of cash transfers in Hungary (van de Walle et al. 1994). A more recent example includes Younger (2001 and 2002) for Peru.

The method has limitations. Simple comparisons of incidence at different points in time do not reveal which factors are responsible for marginal incidence patterns; for example, a key question is often to what degree government policy as opposed to income growth can be credited with improvements in equity. The Malaysia study (Hammer et al., 1995) supplements the incidence analysis with more detailed analysis of the underlying mechanisms. Success in the education sector is attributed to the government's policy of ethnic targeting, while pro-poor improvements in the health sector are due to the private sector's increasing ability to attract wealthier households. Another limitation is that when the underlying population distribution alters between periods — due, for example, to urbanization — the technique is unable to disaggregate incidence results over time geographically which is typically of interest.

In implementing ex-post marginal incidence analysis, an issue arises concerning the definition of the quantiles j. Do we want to know whether the amount going to the relatively or to the absolutely poor has changed? The two could be quite different depending on how absolute poverty is changing. Some studies simply define quantiles specific to each date, so that they are not strictly comparable. This helps answer questions concerning changes in incidence for (say) the bottom 20% of the population at any one date. But often, the interest is in how the amounts received, conditional on real income, have changed. Then one would want to fix the cutoff boundaries in real income space rather than in relative income space. When using a panel to study the incidence of the changes in social spending, households can be ranked by three different definitions of welfare, which can be loosely referred to as delineating the initial, new, and long-term poor — namely the welfare indicator in the initial period, the same in the later period and by the mean over both years (see Table 3.2 for an example of the first two).

3.3.3. Marginal incidence analysis using panel data

Two studies have explored the dynamic marginal impacts of public expenditures using ex-post benefit incidence with data that follow the same households over time. This allows an assessment of dynamic marginal incidence. Using a panel of Hungarian households for 1987 through 1989, Ravallion, van de Walle and Gautam (1995) devise a methodology to examine how well the social safety net protected vulnerable households from falling into poverty versus how well it promoted households out of poverty. The essential idea is to simulate a counter-factual joint distribution of the welfare indicator over time without the change in transfers, using similar econometric methods to those described above.

Similar techniques were used in a study of the safety net in Viet Nam (van de Walle 2002b). Poverty fell quite dramatically in Viet Nam and there was a clear expansion in the total outlays going to social welfare programs between 1993 and 1998. This provides an interesting backdrop for a number of questions concerning marginal incidence: Was the expansion pro-poor? What role did transfer programs play in the reduction in poverty? Does the revealed instability over time in who gets transfers reflect the system's response to changing household circumstances?

An important role for the public sector in a poor rural economy like Viet Nam is to provide protection for those who are vulnerable to poverty due to uninsured shocks. The typical incidence picture is uninformative about whether transfers perform such a safety net function. The static average incidence may not seem particularly well-targeted, but it may be deceptive about the degree to which outlays, coverage, and changes over time, were perhaps correlated to poverty related shocks and changes in exogenous variables.

Table 3.2 ranks households by two definitions of welfare as discussed above—namely per capita expenditures (net of half of transfers) in the initial period and the same in the later period —and presents a comparison of the average incidence of social income receipts in both years and the marginal incidence of the spending increase. The former is expressed as the percent of total social income going to each quintile while the latter is given by the percent of the total increase going to each group. In this particular case, little difference is found between the average and marginal incidences for either definition of welfare. Expansion was more or less proportional to base year receipts across groups. The evidence does not suggest that the poor were specifically targeted by the program expansion.

Table 3.2: Distribution of social transfer income in Viet Nam

	% of 1993 transfers	% of 1998 transfers	% of total transfer increase 1993-98
1993 Net quintile:			
1	13.3	13.1	12.8
2	15.2	15.5	15.7
3	16.9	17.5	17.9
4	21.2	22.4	23.4
5	33.3	31.6	30.2
Total	100.0	100.0	100.0
1998 Net quintile:			
1	15.3	15.7	16.0
2	14.0	15.4	16.6
3	20.5	19.6	18.9
4	19.8	19.9	20.0
5	30.4	29.3	28.5
Total	100.0	100.0	100.0

Source: 1993 and 1998 VLSS.
Note: National population quintiles are constructed using per capita expenditures net of half of social transfers.

Were changes in transfers responsive to poverty-related shocks? Table 3.3 presents information on mean changes in transfers received by panel households classified into a three by three matrix. Households ranked into terciles of their initial 1993 level of per capita consumption (low, middle or high) are cross-tabbed against the change in their consumption between the two dates categorized into whether it underwent a fall, stayed more or less the same or rose significantly. The results strongly suggest that the programs were unresponsive to consumption shocks. Neither starting out poor, nor experiencing negative consumption shocks, appear to have elicited a response from social welfare programs.

The study also examines what role transfers played in the impressive reduction in poverty that occurred over this period. The panel structure is exploited to evaluate how well the safety net performed dynamically following the approach proposed in Ravallion et al. (1995). In comparing joint distributions of consumption expenditures, such as with and without policy changes, the approach tests a policy's ability to protect the poor (PROT) and its ability to promote the poor (PROM). It indicates which distribution offered more protection and which offered more promotion and allows a calculation of the statistical significance of the difference.

Table 3.3: The incidence of changes in transfers by initial consumption and changes in consumption over time

	Fall in consumption	Consumption stayed the same	Large rise in consumption
Low initial consumption	34%	27%	27%
	111,901	246,476	241,658
	80	506	848
Middle initial consumption	32%	30%	30%
	408,469	251,619	296,513
	240	422	772
High initial consumption	33%	36%	32%
	481,618	343,329	367,991
	496	221	720

Source: van de Walle (2002b) using the 1993 and 1998 VLSS.
Note: The population is ranked into three equal groups based on 1993 per capita expenditures net of half of transfers and cross-tabbed against the level of their change in consumption over time net of half the change in transfers. The first number gives the % of households in the cell who received transfers in 1998. The second number gives the per capita amount of the change in transfers received by those with positive receipts only. The final number gives the number of households in the cell. Changes in transfers refer to changes in amounts received from social insurance, social subsidies and school scholarships.

The baseline joint distribution of consumption in the two years is presented in Table 3.4. Households are classified into four groups according to whether they were poor or non-poor in both years, and whether they escaped or fell into poverty over the period. There is evidence of a large fall in poverty and considerable persistent poverty. What was the effect of transfers on poverty? To answer this question, it is necessary to simulate the counterfactual joint distribution without transfers; as in the static incidence calculations, this is done by subtracting half the transfers received in each respective year from consumption in that year. Transfers are found to have negligible impact on poverty. Without them, one and two additional percent of the population would have been poor in 1993 and 1998 respectively.

One can also assess the impact of changes in transfers between the two dates by simulating the joint distribution had there been no changes. Or ask how much better targeting could have improved impacts on poverty incidence by comparing the current distribution relative to a simulated uniform allocation of actual 1998 social income across the entire population. A number of targeting scenarios can be tested.

The Viet Nam analysis concludes that transfers had a negligible bearing on poverty outcomes and failed to protect those who faced falling living standards during this period. By contrast, the Hungary case study found that cash benefits protected many from poverty though it helped few escape poverty. The policy reforms examined were more successful at reducing transient than persistent poverty. Such evidence on dynamic performance is of key value in designing effective safety nets.

Table 3.4: The baseline discrete joint distribution

	1998		
1993	Poor	Non-poor	total
Poor	33.54% (55.78)	26.58% (44.22)	60.12 100
Non-poor	4.84% (12.14)	35.04% (87.86)	39.88 100
Total	38.38	61.62	100

Source: van de Walle (2002b) using the 1993 and 1998 VLSS.
Note: The population is ranked into poor, non-poor groups based on actual per capita expenditures at each date and cross-tabbed. The first number in each cell gives the % of total population who were in that row's poverty group in 1993 and that column's group in 1998. Numbers in parentheses give the proportion of each row's population that is in each column's group in 1998 or the transition probability.

Table 3.5: Joint distribution without transfers

PROT= 0.31(0.66); PROM= 0.70(0.74)

	1998		
1993	Poor	Non-poor	total
Poor	35.21% (57.63)	25.88% (42.37)	61.09 100
Non-poor	5.15% (13.24)	33.76% (86.76)	38.91 100
Total	40.36	59.64	100

Source: van de Walle (2002b) using the 1993 and 1998 VLSS.
Note: The population is ranked into poor, non-poor groups based on their simulated without transfer per capita expenditures (minus .5*transfers) at each date and cross-tabbed. z-scores in parentheses outside the table; critical values: 1.96 (2.58) at the 5% (1%) level.

3.3.4. Marginal incidence analysis using geographic panel data

Another strand of the literature has focused on tracking the incidence of public spending across geographic areas over time. This does not require household level panel data. Instead, the idea is to aggregate cross-sectional data geographically (typically these will be sub-national governmental areas such as provinces) and compare incidence over time. By allowing for geographic fixed effects in incidence, the method can study incidence and its determinants robustly to differences between local governments in preferences for redistribution or other sources of latent heterogeneity. Ravallion (1999) illustrates this approach using province level panel data on the incidence of a social program in Argentina.

The first step is to estimate benefit incidence, or a summary measure of incidence, for each geographic area at each date. For this purpose one might use household survey data or data for even finer geographic areas; for example, Ravallion uses the empirical relationship between program spending and the local poverty map across local government areas ("departments") within each province of Argentina. The poverty map uses the percent of households with unmet basic needs at departmental level based on census data. The spatial variances in both spending and poverty incidence within each province are exploited to measure targeting performance. This entails running an OLS regression for each province of spending allocations across departments against poverty incidence. The regression coefficient gives a "targeting differential" interpretable as the mean difference in spending between the poor and non-poor (Ravallion, 2000). This is done for provinces and all dates for which spending allocations are available. Thus a panel can be formed of these targeting differentials by provinces and over time.

Next the province and date specific targeting differentials are regressed on per capita spending allocations to the provinces in a regression that pools all dates and provinces and includes a province fixed effect to capture province specific factors that affect targeting. This allows one to see what happens to targeting performance during program expansions and cutbacks. Ravallion (1999) finds that during a retrenchement period in Argentina's Trabajar program, a $1 cut in average spending reduced the targeting differential by $3.55 on average. Hence, cuts were accompanied by worsening targeting performance.

3.4 Conclusions

Ignoring behavioral responses to public spending or social programs can yield deceptive assessments of incidence because one does not correctly assign beneficiaries to the pre-intervention distribution. For example, by subtracting the full amount of a transfer received one overestimates how poor beneficiaries will have been in the absence of the intervention. Similarly, if one ignores how the political economy can influence the assignment of beneficiaries and the incidence of program spending allocations one can arrive at severely biased estimates of impacts and hence incidence.

This chapter has reviewed some relatively simple methods that can help address these deficiencies of non-behavioral incidence analysis. There is a large literature on reduced-form regression-based methods in which one essentially regresses the relevant outcome measure (income for example) on program allocations with relevant controls. With household panel data one can exploit changes in program spending over time to obtain estimates that are robust to potential endogeneity of the program assignment across units (provided the endogeneity is fully reflected in time invariant factors). These methods can be made more sophisticated, such as by incorporating heterogeneity in impacts and allowing for more complex forms of endogeneity in program assignments or spending levels. Once one knows the impact on incomes one can then work out what the income or other welfare indicator of program participants would have been in the absence of the intervention, and so estimate the incidence of spending relative to that counter factual.

Marginal benefit incidence analysis is another important example of how behavioral responses through the political economy of incidence can be incorporated. This can provide valuable information for charting the course of pro-poor reforms in public spending. The method can be implemented using a single cross-sectional survey, but access to two or more consecutive household cross-sectional surveys with information on program participation will usually provide estimates that are more robust. Household or regional panel data allow an even richer analysis of policy and spending changes. By examining actual changes ex post, these methods provide a reality check for the results of methods that attempt to approximate or predict reality ex ante.

3.5. References

Jalan, Jyotsna and Martin Ravallion, 2002, "Estimating the Benefit Incidence of an Antipoverty Program by Propensity Score Matching," Journal of Business & Economic Statistics, forthcoming.

Hammer, Jeffrey, Ijaz Nabi and James Cercone, 1995, "Distributional impact of social sector expenditures in Malaysia," in Public Spending and the Poor: Theory and Evidence, (eds.) D. van de Walle and K. Nead, pp. 521-554. The Johns Hopkins University Press, Baltimore and London.

Lanjouw, Peter and Martin Ravallion. 1999. "Benefit Incidence and the Timing of Program Capture," World Bank Economic Review, 13(2): 257-274.

Ravallion, Martin, 1999, "Is More Targeting Consistent with Less Spending?" International Tax and Public Finance 6: 411-19.

Ravallion, Martin, 2000, "Monitoring Targeting Performance When Decentralized Allocations to the Poor are Unobserved," World Bank Economic Review 14(2): 331-45.

Ravallion, Martin, Dominique van de Walle and Madhur Gautam, 1995, "Testing a Social Safety Net," Journal of Public Economics 57: 175-199.

van de Walle, Dominique, 2002a, "Poverty and Transfers in Yemen," Middle East and North Africa Working Paper Series No. 30, World Bank, Washington, DC.

van de Walle, Dominique, 2002b, "The Static and Dynamic Incidence of Viet Nam's Public Safety Net," Policy Research Working Paper No. 2791 Development Research Group, World Bank, Washington, D.C., February 2002.

van de Walle, Dominique, 2002c, "Viet Nam's Safety Net: Protection and Promotion from Poverty?," mimeo, Development Research Group, World Bank, Washington, D.C.

van de Walle, Dominique, 1994, "The Distribution of Subsidies through Public Health Services in Indonesia 1978-87," World Bank Economic Review 8(2): 279-309.

van de Walle, Dominique, Martin Ravallion and Madhur Gautam, 1994, "How Well Does the Social Safety Net Work? The Incidence of Cash Benefits in Hungary 1987-89," Living Standards Measurement Study Working Papers, 102.

Younger, Stephen, 2002, "Public Social Sector Expenditures and Poverty in Peru," in Christian Morrisson, ed., Education and Health Expenditure, and Development: The Cases of Indonesia and Peru, OECD, Development Center Studies, Paris.

Younger, Stephen, 2001, "Benefits on the Margin," mimeo, Department of Economics, Cornell University.