BENEFITS ON THE MARGIN:

OBSERVATIONS ON MARGINAL BENEFIT INCIDENCE*

Stephen D. Younger

Cornell University
Food and Nutrition Policy Program
Room 3M12 Martha van Rensselaer Hall
sdy1@cornell.edu

October, 2002

*I have benefited from the comments of Peter Glick, Francois Bourguignon, participants in the OECD research program "Development of Human Resources and Poverty Alleviation," and two anonymous referees. Any remaining errors are my own. Part of this research is supported by USAID under cooperative agreement AOT-0546-A-00-3178-00.

ABSTRACT

Benefit incidence analysis has become a popular tool over the past decade, especially for researchers at the World Bank (Demery 1997, van de Walle and Nead 1995, Selden and Wasylenko 1992). Despite, or perhaps because of, the popularity of this method, more recent research has pointed out many of its limitations (van de Walle 1998, Lanjouw and Ravallion 1999). One of the most common criticisms of the standard benefit incidence method is that its description of *average* participation rates is not necessarily useful in guiding *marginal* changes in public expenditure policy from the status quo.

This paper considers a variety of options for analyzing the marginal benefit incidence of policy changes. A key conceptual point is that, despite the fact that each method measures "marginal" incidence, they do not in fact measure the same thing, nor are they intended to do so. There are many possible policy changes, and thus many margins of interest. Each method captures one of these, and so is at least potentially of interest for some analyses, while potentially inappropriate for others. Empirically, the precision of the methods differs substantially, with those relying on differenced data or aggregations of households into groups yielding standard errors that are quite large relative to the estimated shares.

I. INTRODUCTION

The past decade has witnessed a resurgence of interest in the relation between public expenditure and poverty in developing countries. This resurgence has fostered the return of incidence analysis, particularly with respect to the benefits of public expenditures in the social sectors. While analysis of tax incidence has a long and venerable history in economics, distributional analysis of the benefits of public expenditures (or public policy more generally) is more recent (Aaron and McGuire 1970, Brennan 1976, Meerman 1979 Selowsky, 1979). Broadly stated, this so-called benefit incidence analysis addresses the question "How are the benefits of government expenditures on X distributed across the population?"

While there are potentially many ways to approach this question, a fairly standard method has emerged, largely from researchers at the World Bank (Demery 1997, van de Walle and Nead 1995, Selden and Wasylenko 1992). This method takes "across the population" to be "across the expenditure (or income) distribution," which is consistent with the overall concern with poverty. It then uses some variant of an average participation rate in a public program for people in different strata of the expenditure distribution to estimate the distribution of benefits. Given a presumed preference for public expenditures that benefit the poor, programs or policies in which the poor have higher average participation are viewed more favorably. The large increase in availability of nationally representative, multi-purpose surveys such as the World Bank's Living Standards Measurement Survey (Grosh and Glewwe 1998) and the relative ease with which this standard method can be applied have led to a profusion of such analyses. It is now common to find a benefit incidence analysis in any developing country's poverty profile and in many project proposals and evaluations.

Despite, or perhaps because of, the popularity of this method, more recent research has pointed out many of its limitations (van de Walle 1998, Lanjouw and Ravallion 1999). One of the most common criticisms of the standard benefit incidence method is that its description of *average* participation rates is not necessarily useful in guiding *marginal* changes in public expenditure policy from the status quo, a point first made by Lipton and Ravallion (1995). The logic of this argument is compelling. The standard method describes who is benefiting from a particular public expenditure now. As such, it is a useful guide to the consequences of a policy change whose benefits are distributed in proportion to current benefits. But there is no necessary reason that a policy to increase expenditures will go to existing beneficiaries in proportion to their current benefits, or that it will go to existing beneficiaries at all. Many policies explicitly aim to expand the benefits of public expenditure to non-beneficiaries. In this case, since the benefits by definition do not go to existing beneficiaries, the standard method is misleading. Even when one envisions a change in the characteristics of existing services, the changes may not be uniform across existing users, in which case the standard method is also inappropriate. The benefits of a policy to ensure that all students have a complete set of textbooks, for example, will have different distributional consequences if some students already have a complete set, and so gain nothing, while others do not.

In response to this observation, several recent papers have proposed alternative methods to measure the marginal incidence of public expenditures. Glick and Razakamanantsoa (2001) and Younger (2002) look at shares of the change over time in the share of benefits across the expenditure distribution. Lanjouw and Ravallion (1999) and Galasso and Ravallion (2001) estimate the "marginal odds of participation" for each expenditure quintile as the coefficient in a regression of quintile/small area participation rates on large area participation rates. Lanjouw, and others (2002) and Ravallion (1999) apply similar techniques to panel data in order to control for fixed area characteristics. Younger (1999, 2002) considers marginal incidence to be the distribution of compensating variations for marginal policy changes, based on estimated demands for public services.

This paper considers each of these options for analyzing the marginal benefit incidence of policy changes using a specific example of secondary education in rural Peru. A key conceptual point is that, despite the fact that each method measures "marginal" incidence, they do not in fact measure the same thing, nor are they intended to do so. There are many possible policy changes, and thus many margins of interest. Each method captures one of these, and so is at least potentially of interest for some analyses, while potentially inappropriate for others. Empirically, the precision of the methods differs substantially, with those relying on differenced data or aggregations of households into groups yielding standard errors that are quite large relative to the estimated shares. This result argues for caution with these methods when using samples similar to the Peru surveys, which in turn are about the same size as many existing multipurpose household surveys.

II. METHODS

The Standard Benefit Incidence Method

A standard benefit incidence study requires two main components: a measure of the value of the benefit that an individual, household, or population sub-group receives from a particular public expenditure; and a way to compare the beneficiaries to the population in general. When studying the benefits of public services, the standard method usually uses the government's cost of provision to estimate the service's value to users. But there are both theoretical and practical reasons to doubt this practice (van de Walle 1998, Sahn and Younger 2000), so an increasing number of evaluations simply count users, i.e. a user or beneficiary gets a benefit of one, others get zero.

It is possible to compare the beneficiaries of a public expenditure to the general population along many dimensions - for example ethnicity, gender, region of residence, age, functional income classifications, or political constituency – but an interest in poverty and inequality implies that most of our comparisons will involve welfare. That is, we want to know how the recipients' welfare compares to the general population's welfare. Almost all work for developing economies uses household expenditures per capita or per adult equivalent as its measure of welfare. Once we have decided how to value benefits and how to group the sample, the calculations are simple: divide each individual's or household's benefit by the total to get his/her share of benefits, and sum those shares across a population sub-group, usually welfare quantiles.

This standard method clearly uses sub-group averages to estimate the distribution of benefits. Nevertheless, this average measure does yield the distributional consequences of a marginal policy change whose benefits are distributed to existing users in proportion to their benefit. The most obvious of these would be a tax or subsidy that changes the existing price proportionately, but one can think of others, such as a new uniform for each child in school or a new vaccination made available to each member of a social security system. As such, the terminology that compares "average" benefit incidence, calculated in the standard way, to "marginal" benefit incidence, calculated with one of the other methods mentioned in the introduction, is unfortunate. The standard method does capture a margin and can be interpreted as such in terms of welfare theory (Yitzhaki and Slemrod 1991).

Rather, the problem with the standard method is that this is not the margin that interests most people. For example, policy makers often do not think of increases in public expenditure for health or education in terms of larger price subsidies for those services. Instead, they have in mind an expansion of these services to non-beneficiaries induced by increased access rather than a reduction in price. As I have noted, such benefits do not go, by definition, to existing beneficiaries, so the standard method is inappropriate. In the next section, I discuss several methods to estimate non-proportional expansions of public service coverage.

Estimating the Benefits of a Marginal Expansion of Services

METHOD 1A: Using spatial variation in coverage to estimate marginal program benefits. Lanjouw and Ravallion (1999) develop a political economy model in which the share of discrete population groups, e.g., the poor and non-poor, have different political power, and different costs and benefits from a given public expenditure.[1] The interplay between these factors determine the relationship between program size, the total public expenditure on a program or service, and the share of each group in that program's benefits.[2] "Early capture" by the poor occurs when they receive higher shares of small programs, but their share declines as the program size increases. "Late capture" is defined as the opposite case.

Even with substantial restrictions, the theoretical model yields no general results as to whether early or late capture will occur, so the question is empirical. To address it, Lanjouw and Ravallion estimate the following regression:

$$p_{j,k,q} = \alpha_q + \beta_q p_k + u_q \qquad (1)$$

where j indexes a small geopolitical unit (a province in Peru), k indexes a larger one (a department in Peru), and q indexes the welfare quantile. The left-hand variable is the program participation rate for a given province and quantile. The regressor is the program participation rate for the department in which that province is located. β_q, then, is the marginal effect of an increase in program participation for the entire department on the participation rates for people in a given province and quintile. We run the regression for each quintile separately. Also, because $p_{j,k,q}$ is included in p_k, there is an upward bias in the estimate of β_q. Lanjouw and Ravallion

[1] Similar models are found in Ravallion (1999, 2002).
[2] In the particular specification of Lanjouw and Ravallion, the non-poor bear all the program costs, and they also hold all of the political power in the sense that the poor cannot impose on them a program that lowers their welfare. In such a case, convexity of the program cost function is sufficient to guarantee "early capture" by the poor.

resolve this by instrumenting p_k with the left-out mean, that being the participation rate for all of department k except those individuals in province j and quantile q.

The intuition behind the regression is that, by observing variation in departmental participation across the country, we can understand how greater coverage affects the participation of sub-groups. If β_q is greater than one, it indicates that a general expansion in coverage is correlated with a disproportionately large increase in participation for that province/quintile. One advantage of this method is that it requires only a cross-section of data, just as the standard method. An important assumption is that the political process that determines the correlation between program size or coverage and incidence is the same across regions.

The margin that this model estimates is the incidence of an increase in program participation. The model does not address which policy or policies might bring about the program expansion, nor does it consider specific changes in demand for services. Rather, it makes a more general appeal to the political economy behind the policies to argue that, whatever the specific policy used – price reductions, quality improvements, or reduced rationing – the outcome must respect the political constraints that each group's costs, benefits, and political power imply.

METHOD 1B: Controlling for fixed effects. Lanjouw and Ravallion point out that equation (1) has no controls for any effect on quintile/province participation rates except the department's participation rate. In cases where surveys are available for more than one point in time, it is possible to construct a panel of provinces, and thus to include a province fixed effect to control for left-out covariates that are constant over time.[3] This is possible even if the survey is not a panel of households, as long as the households are sampled from the same provinces and each survey is representative at the province level.[4] I compare this variation to the original Lanjouw/Ravallion method in the example that follows.

METHOD 1C: Using disaggregated (individual) data. A purely statistical problem with the Lanjouw/Ravallion model is that it uses province/quantile average data. While this was often a necessity in the past, it is now standard to have access to household-level datasets with which to estimate benefit incidence. Grouping observations into province/quantile averages reduces the efficiency of the estimates, yielding larger estimated standard errors (Johnston, 1972). In the application below, I estimate the model on both group averages and household data.

METHOD 2: Observing changes as programs expand over time. This method addresses the same margin as Method 1: what is the incidence of increased expenditures as a program or service expands? But rather than using the spatial variation in the correlation between program size and incidence, this method calculates each group's share of observed changes in benefits. As such, it is mechanically similar to the standard benefit incidence method, except that it

[3] Lanjouw and others (2002) and Ravallion (1999) pursue this strategy.
[4] The first condition is often true of household surveys, while the second is quite rare. In the two Peru surveys that I use in the next section, there is considerable overlap of provinces (the small geopolitical unit), but the samples are not representative at the province level, leading to the possibility that the observed variation over time is due to sampling differences. Fortunately, the ENNIV usually returns to the same clusters when it conducts a new survey, which should minimize this problem.

substitutes the change in a given quantile's benefits (or program participation) for its level. Glick and Razakamanantsoa (2001) and Younger (2002) uses two cross-sections at different points in time to estimate each quantile's share in the *change* in use of various public services.[5]

This method requires at least two cross-sectional surveys, but just as the number of developing countries with at least one nationally representative multipurpose survey grew in the previous decade, an increasing number of countries now have more than one such survey a few years apart. Like Method 1, this approach also says nothing about the incidence of program expansions brought about by particular policy instruments. It is purely a description of what actually took place between the two surveys in terms of program coverage and shares.

<u>METHOD 3A: Econometric estimates of compensating variations</u>. Rather than use the standard benefit incidence method as an approximation to the compensating variation to a price change, it is possible to estimate compensating variations, for price and other policy changes, econometrically. There is a well-established literature in transport and environmental economics that does this for goods and services where demand is discrete (Small and Rosen, 1981, McFadden, 1995) and Gertler and his associates have applied the techniques to health and education demand in developing countries (Gertler, Locay, and Sanderson, 1987, Gertler and Glewwe, 1990).

The model is well known. We assume that each household has a utility function that depends on its consumption and on the quality of the school choice that it makes:

$$V_j = f[y-p_j,Q(X_j,Z)] + e_j \qquad (2)$$

where j indexes the choice (no school or school); y is household permanent income, proxied by household expenditures; p_j is the price of choice j, including all opportunity costs of time; Q is a function that measures quality, which depends on choice-specific characteristics X_j and on household or personal characteristics Z. The household chooses the option j that yields the highest utility. Even though V_j is not observable, we know that if a household chooses option j, V_j is greater than all other V_i. The model estimates the probability that this is so, using only the observed choice, and takes the probability of choosing option j as an expected demand for that option. Small and Rosen (1981), then, show how to calculate compensating variations in such a model.

An important identifying assumption of a model like equation (2) is that the observed choice is actually the one that provides the highest utility, which implies that there is no rationing beyond what can be captured with the choice-specific characteristics, X_j. For many public services, this may not be the case. For example, schools may exclude students based on merit, gender, social status or connections, etc. In such cases, we may observe students who are not attending school even though that is the option that provides them the greatest utility, which

[5] Note that Van de Walle (1995), Hammer, Nabi, and Cercone (1995), and Lanjouw and others (2002) also use two cross sections, but they describe how the standard benefit incidence changes over time rather than the incidence of the changes.

would bias the estimates in equation (2) if the rationing is correlated with the regressors.[6] In the specific case study that follows, rationing is not a problem. With the exception of a few prestigious schools in urban areas, secondary schools in Peru do not ration slots. This is possible because class sizes are not limited.

Unlike methods 1 and 2, this method permits traditional policy analysis in the sense that it answers the question "Who will receive the marginal benefits if we change policy variable x_j?" As such, it is useful for estimating the marginal incidence of any policy change for which appropriate x_j data are available.

METHOD 3B: Econometric estimates of changes in the probability of participation. The method based on compensating variations differs from the other methods presented in that it considers the value of a policy change to potential recipients rather than the change in (probability of) participation. This valuation adds an extra dimension not found in the other methods. Glick and Sahn (2000) estimate the model in equation (2), but calculate only the change in the probability of participation associated with simulated policy changes. By modeling participation rather than its monetary value, this approach is closer to the others presented than method 3a. An advantage of this method over the estimation of compensating variations is that, because it models only the probability of a given option, it remains valid in the presence of rationing.

III. BENEFIT INCIDENCE OF SECONDARY SCHOOLING IN PERU

In this section, I calculate estimates of the distribution of benefits from an expansion of secondary school attendance in Peru, using all of the methods outlined in the previous section. In addition, I calculate the standard benefit incidence measures. The calculations for the shares in observed changes in participation, Method 2, are straightforward. We only need to note that overall coverage for secondary school increased slightly, from 6.33 to 6.58 percent of the rural population, between 1994 and 1997. The other methods require preliminary regressions, which I report first.

Table 1 presents the results of four Lanjouw/Ravallion-type regressions described as Method 1. The data for the cross-section are for rural households from the 1994 and 1997 rounds of Peru's *Encuesta Nacional de Hogares sobre Medición de Niveles de Vida* (ENNIV). The regressions are for rural provinces only, to be consistent with the model for method 3 below. In the upper right quadrant of the table, the dependent variable is a 0/1 indicator of household-level participation in secondary schooling.[7] In all of the other quadrants, the data for the dependent variable are the province- and quintile-specific participation rates for attendance at secondary school, defined as the number of secondary students divided by the population. In all models, the right-hand variable is the department-wide participation rate, and all estimates are two-stage least squares, using the left-out mean department participation rate as an instrument.

---

[6] Some non-price rationing that is characteristic of the *service* can, however, be modeled. For example, health centers may charge low fees and handle the excess demand by imposing long waiting times. As long as the waiting time can be included in X_j, the estimates are consistent.

[7] So I am estimating the linear probability model. It would also be possible to estimate this model as a probit or logit.

In addition, I have imposed two restrictions on the coefficients: that the α_q's sum to zero and that the β_q's sum to the number of quantiles, five in this case. While Lanjouw and Ravallion do not impose this restriction, they are required if the estimated shares of marginal benefits are to sum to one. As it happens, the unrestricted estimates are quite close to those reported here for all the cross-section models. For the panel data model, the differences are much larger, but so are the standard errors, so that even for this model the slope coefficients do not differ at the five percent level.

While I will postpone a discussion of the distributional implications until the following section, there are three points to note about the results in Table 1. First, the various models produce quite different estimates for the quintile-specific marginal odds of participation. This is true even for the two cross-section models, in the upper and lower left-hand quadrants. These differences cannot be due to changing marginal odds of participation as coverage expands because rural secondary school enrolments were virtually constant between 1994 and 1997 in these samples. While the standard errors are large, there are several significant differences, especially for the poorest quintile's marginal share. Second, in all the models, the poorer quintiles receive a less-than-proportionate share of marginal benefits from secondary schooling. Finally, as expected, the standard errors for the slope coefficients are only about half as large for the individual level model. Nevertheless, because of the low shares of the poorest quintile, all models reject the null of equal participation across quintiles.

Table 2 gives estimates of the demand for secondary schooling in rural Peru, using the same 1994 ENNIV dataset.[8] While it is customary to consider multiple options for schooling – no school/public school/private school, or no school/local school/distant school – in rural Peru, only 3 percent of children attend private schools, making an estimate of demand for the private options infeasible, and the survey does not identify students who are away from home studying. Thus, I estimate only the choice between attending school or not, as a probit. I limit my attention to rural areas because a model with only a few choices is not appropriate for most urban areas. A resident of Lima has a choice of many schools, public or private. No survey in Peru permits us to adequately identify, let alone model, these choices.

Since the probit can identify the model only up to the differences in V_j, we must normalize against one option, which will be the no school choice here. Thus, I assume that $Q(X_0,Z)=0$. In my estimates, I assume that the function f() in equation (2) is quadratic in net expenditures, and that it is constrained to be the same for each option. There is some debate about this latter restriction in the literature (Dow, 1999), but as Gertler and Glewwe (1990) note, it is necessary to get a sensible estimate of the marginal utility of income which, in turn, is necessary to calculate compensating variations (Small and Rosen, 1981). The function Q() is linear and separable from net expenditures, except for an interaction between net expenditures and distance to school. This variable is the one that I will use to compare results to the Lanjouw/Ravallion method, so I want to allow as much flexibility as possible.[9]

[8] I use the 1994 data because they include a richer set of questions on school quality, including distance to school and parents' evaluation of problems at their children's schools.
[9] As it turns out, the estimated changes in probability of attendance and compensating variations from this model have correlation coefficients greater than 0.95 with those from a model without the interaction.

The samples include all rural children who are either attending school or who are eligible to attend school. The latter group includes all children of the appropriate age who have not yet graduated from secondary school. I include even children of secondary age who have not graduated from primary school, the argument being that, in the context of a long-run optimization, the decision not to complete primary school is affected in part by perceptions of the value of secondary school. Dow (1999) defends this type of unconditional estimate.

All of the household characteristic variables in Table 2 are self-explanatory except for net household expenditures. Household expenditures are defined in the broadest way possible, including imputed value of owner-occupied housing and own-produced goods (Younger, 2002). In addition, if the household has a student, I add back the costs of schooling and the opportunity cost of time at school to get a gross expenditure variable that is before the costs of schooling. The price is the cluster- or district-level mean for school costs, including fees, books, uniforms, transportation, and the opportunity cost of time.[10] If the cluster had at least four observations, I used only the cluster-level data to calculate means. Otherwise, I used district-level data.[11] All means calculated in this way are left-out means. Net expenditures are gross expenditures minus the price.

The child characteristics are also self-explanatory. The default option for relation to the household head is being the head him/herself. The school characteristics require more explanation. The number of required books is the cluster- or district-level mean number of textbooks that the school requires. I take this as an indicator of higher academic quality. Distance to the school is in kilometers and time in minutes.

The questions about parents' wish to change features of their school are based on the following question to one adult per household, if the household includes children: "If you could change anything about your children's school, what changes would you make? (Use a scale from 1 to 3)." This is followed by the list of features used as regressors. I have not used the scale, but simply recorded whether the parent expressed a desire to improve a feature. I then use the cluster- or district-level share of parents expressing in interest in improving each feature in the regression.

The signs of the coefficients are almost all consistent with one's prior expectations. Net household expenditures have a positive, and only slightly concave, effect on the probability of secondary school attendance in rural areas. Children in households whose head is older, urban-born, and more educated are more likely to attend. Children in households with younger children (ages zero to twelve) are less likely to attend, while those in households with older children (exclusive of the child whose observation this is) or adults are more likely to do so. Of the child's characteristics, only one variable is statistically significant at standard levels: girls who are married are less likely to attend secondary school, even after accounting for the positive effect coming from being the household head's spouse (almost all of whom are women). Being female and coming from a household with at least one indigenous language speaker also lowers the probability of attendance, but the t-statistics are smaller for these variables. All but one of the school characteristics has significant effects on the probability of attending. On the other

[10] Districts are the third-level geopolitical units in Peru, smaller than provinces.
[11] I use the same criterion of at least four observations for all of the subsequent cluster- or district-level regressors.

hand, only one of the parental opinion questions is significantly different from zero, that expressing a desire for "improved desks and services."

I use the results in Table 2 to simulate two policy changes. The first is a reduction in fees of 100 soles, about the sample average expenditure per student on school fees, books, uniforms, and transportation. This is a policy change that the standard benefit incidence method approximates, so the distribution of estimated benefits from the two methods should be close. The policy simulation reduces each student's distance to a secondary school to a maximum of 2 kilometers. This affects about two-thirds of the sample. School placement is clearly a policy variable, and I believe that it is the variable that many people have in mind when they think of a policy to expand access to public schools. As such, it is the variable in the demand function that is most consistent with Methods 1 and 2. For each simulation, I calculate for each child the compensating variation for the policy change (method 3a), and the change in the probability that s/he attends school (method 3b).

McFadden (1995) shows that the standard method for calculating compensating variations in a discrete choice model developed by Small and Rosen (1981) will yield biased estimates of the compensating variation if utility is a non-linear function of income, as it is in this model. Therefore, I use the simulation method described in McFadden with 1000 repetitions for each observation to calculate the CVs.

Results

Table 3 presents the quintile shares of marginal benefits estimated by each method presented above, with their standard errors.[12] The shares for the Lanjouw/Ravallion methods are those presented in Table 1 divided by five.

Consider first the estimates of the distribution of marginal benefits associated with a program expansion, methods 1a, 1b, 1c, and 2 in columns (A) through (D) of Table 3. Table 1 has already shown that the results for all the Lanjouw/Ravallion methods are statistically different from equal shares, with a modest anti-poor bias.[13] Only a few of the estimates differ by economically important amounts, but with the exception of the difference at the first quintile for methods 1a and 1c, these differences are not statistically significant due to the relatively large standard errors. This is especially true for the panel data method 1b which estimates a much larger marginal share for the top quintile than the other methods, but with a standard error so large that the estimate is not distinguishable from either zero or one.

The results for Method 2 (column D) show an extremely progressive distribution of marginal shares, with the first two quintiles capturing more than 100 percent of the change in benefits. This result that is possible because a quintile can have negative marginal benefits, i.e. a

[12] The quintiles are based on the rural sample of the ENNIV only. I could just as easily derive them for all households in the sample and give zero benefits to urban residents. Where rural residents fall in the nationwide expenditure distribution would then influence the estimated shares — in particular, subsidies to rural secondary schools would look more progressive because rural residents are poorer than urban residents in Peru — but comparisons across methods of each quantile's share would not change.
[13] All tests are at the 5 percent confidence level.

decline in its participation across periods, even as overall participation rates increase.[14] Method 2 seems unsatisfactory on two counts. The resulting marginal share estimates are very different from all of the Method 1 estimates, and also quite erratic. The latter phenomenon can be explained by the fact that the denominator of the marginal shares – the change in attendance at secondary school between 1994 and 1997 – is very small, only 0.25 percent of the rural population. With such a small overall change, any quantile's share of that change can be large, even if its participation did not change much. That said, the method is meant to capture shares of marginal changes which are, by definition, small. Thus, applying it to services with larger expansions might produce more stable estimates, but ones that would be less accurately termed "marginal."

A more important problem with method 2 is that, because it relies on differenced data, the estimates have very large standard errors, so large that the marginal share estimates for this method are statistically indistinguishable from the others despite its very different point estimates. While this is only one example, it seems that Method 2 may require samples that are much larger than is typical in developing countries to produce precise estimates of marginal shares.

I have argued that the policy experiment that most people have in mind when considering a program expansion is "increased access." For secondary schools in rural Peru, this is best captured by reduced distance to school, a variable that is relevant in Peru during this time. The government of Peru invested considerable resources in building and rehabilitating schools in the 1990s through FONCODES (Paxson and Schady 1999), and median travel time to get to school in rural areas declined from 40 to 30 minutes between 1994 and 1997.[15] So even though Methods 1 and 2 apply to *all* changes that affect program size, reduction in distance to school should have been an important factor in this period. As such, it is interesting to compare these methods to Method 3

.

The precise policy change that I simulate is a reduction in distance to school to a maximum of two kilometers.[16] The estimated distribution of the change in the probability of attendance (column G) are somewhat more progressive than that for the compensating variations (column E), though only the fifth quintile's shares differ significantly. This indicates that the value that people place on secondary schooling increases with household expenditures per capita in rural Peru. Both estimates of marginal shares are somewhat larger than any of the Lanjouw/Ravallion methods for the poorest quintile. For the changes in probability (column G), all of these differences are statistically significant, but only the first quintile difference with method 1c, the more precise of the Lanjouw/Ravallion methods, is statistically significant for the compensating variations (column E). Method 3b also differs significantly from Method 1c (the individual level model) at the third and fifth quintiles, and from Method 1a (the province-level cross-section) at the fifth quintile. For Method 3a, the only other significant difference is for the

[14] This was true for the third and fifth quintiles between 1994 and 1997 in rural Peru.
[15] The 1997 ENNIV does not ask for distance to school, only the travel time.
[16] The mean and median distance in the sample are 4.4 and 2.7 kilometers, with a standard deviation of 7 kilometers. The extremum for distance is 36 kilometers.

shares for the third quintile when compared to Method 1c. The fact that the estimates differ for some quintiles is evidence that more is at work than the distance to school across space and across samples in Peru, which is not surprising.

Finally, quintile shares estimated with the standard benefit incidence model (column I) based on school attendance are quite close to those derived from both the change in probability of attendance (column H) and the compensating variation (column F) for a 100 sole price change, although the difference at the first quintile is statistically significant, if minor (0.03 and 0.05, respectively) for Models 3a and 3b, as is the difference at the fifth quintile for Model 3b. Thus, the standard method yields a good approximation to the marginal incidence of a price change. It is also interesting to note that, unlike the examples cited in Ravallion (2002), including the original Lanjouw and Ravallion result, the shares from the standard benefit incidence method for rural secondary school attendance in Peru are not significantly less pro-poor than those from any of the other methods, except for the simulation of reduced distance to school.

IV. CONCLUSIONS

Benefit incidence analysis is now quite common, partly because of the importance of the issue that it addresses, and partly because it is easy to do. Nevertheless, critics have pointed out that the standard method used to carry out the analysis can often be misleading because it uses quantile average shares of benefits, while analysis of any policy change should be done on the margin. In this paper, I argue that the standard method can in fact be interpreted as a marginal method: it gives a first-order approximation to the distributional consequences of a price change, or any other change that affects only observed beneficiaries, proportionately. In the example of secondary school attendance in rural Peru, I find that the approximation is reasonably good, even for a large (non-marginal) change in the cost of attendance. This is consistent with previous work for five different social services in Ecuador (Younger 1999). In this sense, the paper supports the standard method, as long as it is interpreted correctly.

More broadly, however, there more margins of interest than price. In particular, expanded access to services, rather than changes in fees, is often what policy makers have in mind when considering increased expenditures on a public service. The paper explores different methods that apply to different margins. Methods 1 and 2 do not identify specific causes of a program expansion, but rather argue that however a program expands, it will have to respect political economy constraints which can be captured in the correlation between a program's size and its distribution of benefits. Method 3, on the other hand, is grounded in more traditional policy analysis, identifying the marginal incidence of a specific policy change based on household's compensating variations, or willingness to pay, for that change or, more narrowly, on its probability of participation.

In general, the different methods do produce different estimates of marginal benefit incidence, suggesting that analysts should tailor their choice of method to the question at hand. Those interested in a general description of program beneficiaries or in the incidence of change in benefits proportional to existing use can use the standard benefit incidence method (Method 4) or the methods based on demand estimates (Method 3). Method 2 is appropriate for those interested in a description of the incidence of changes in benefits over time. For the distributional consequences of a general expansion of program coverage in an unchanging political-economic environment, one of the Lanjouw-Ravallion methods (Method 1) is the relevant option. And those interested in analyzing specific policy changes whose impact is not proportional to existing demand should choose the methods based on demand analysis (Method 3).

Apart from these conceptual differences in methods, two important results from the examples relate to the precision of each method's estimates. First, those that rely on individual or household level data yield smaller standard errors than those that use regional aggregations. Thus, a straightforward modification of the Lanjouw-Ravallion method using individual-level is preferable where such data are available. Second, the methods that rely on differencing across time have particularly large standard errors. While not surprising, this makes it difficult to use Method 2 with existing surveys, few of which have enough observations to provide adequate precision.

REFERENCES

Aaron, Henry, and Martin C. McGuire, 1970. "Public Goods and Income Distribution." *Econometrica* 38(6):907-20.

Brennan, Geoffrey. 1976. "The Distributional Implications of Public Goods." *Econometrica*, 44(2):391-99.

Demery, Lionel. 1997. "Benefit Incidence Analysis." Poverty Anchor, Poverty Reduction and Economic Management Network, The World Bank (unprocessed mimeo), January.

Dow, William. 1999. "Flexible Discrete Choice Demand Models Consistent with Utility Maximization: An Application to Health Care Demand." *American Journal of Agricultural Economics* 81(3):680-85.

Galasso, Emanuela, and Martin Ravallion. 2001. "Decentralized Targeting of an Anti-Poverty Program." Draft, Development Research Group, World Bank.

Gertler, Paul, and Paul Glewwe. 1990. "The Willingness to Pay for Education in Developing Countries: Evidence from Rural Peru." *Journal of Public Economics* 42(3):251-75.

Gertler, Paul, Luís Locay, and Warren Sanderson. 1987. "Are User Fees Regressive? The Welfare Implications of Health Care Financing Proposals in Peru." *Journal of Econometrics*. 33(1/2):67-88.

Glick, Peter, and Mamisoa Razakamanantsoa. 2001. "The Distribution of Social Services in Madagascar, 1993-99." CFNPP Working Paper #128. Cornell Food and Nutrition Policy Program, Ithaca, NY.

Glick, Peter, and David Sahn, 2000, "The Demand for Primary Schooling in Rural Madagascar: Price, Quality, and the Choice Between Public and Private Providers." CFNPP Working Paper #113. Cornell Food and Nutrition Policy Program, Ithaca, NY.

Grosh, Margaret E., and Paul Glewwe, 1998. "Data Watch: The World Bank's Living Standards Measurement Study Household Surveys." *Journal of Economic Perspectives* 12(1):187-96.

Hammer, Jeffrey, Ijaz Nabi, and James Cercone. 1995. "Distributional Effects of Social Sector Expenditures in Malaysia, 1974 to 1989." In Van de Walle and Nead, op.cit.

Hausman, Jerry, ed. 1993. *Contingent Valuation: A Critical Assessment*. Amsterdam: North-Holland.

Johnston, J.. 1972. *Econometric Methods*. New York: McGraw Hill.

Lanjouw, Peter, and Martin Ravallion. 1999. “Benefit Incidence, Public Spending Reforms, and the Timing of Program Capture.” *The World Bank Economic Review*. 13(2):257-73.

Lanjouw, Peter, Menno Pradhan, Fadia Saadah, Haneen Sayed, and Robert Sparrow. 2002. “Poverty, Education, and Health in Indonesia: Who Benefits from Public Spending?” In Christian Morrisson, ed., *Education and Health Expenditure and Development: The Cases of Indonesia and Peru*. Paris: OECD.

Lipton, Michael, and Martin Ravallion. 1995. “Poverty and Policy.” In Jere Behrman and T.N. Srinivasan, eds., *The Handbook of Development Economics*, volume 3B: 2551-2657. Amsterdam: North Holland.

McFadden, Daniel. 1995. “Computing Willingness-to-Pay in Random Utility Models.” mimeograph, Department of Economics, University of California, Berkeley.

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

Paxson, Christina, and Norbert Schady. 1999. “Do School Facilities Matter? The Case of the Peruvian Social Fund (FONCODES).” World Bank Policy Research Working Paper No. 2229.

Ravallion, Martin. 1999. “Is More Targeting Consistent with Less Spending?” *International Tax and Public Finance* 6(3):411-19.

Ravallion, Martin. 2002. “Who is Protected? On the Incidence of Fiscal Adjustment.” Paper presented at the IMF Conference on Macroeconomic Policies and Poverty Reduction.

Sahn, David, and Stephen Younger. 2000. “Expenditure Incidence in Africa: Microeconomic Evidence.” *Fiscal Studies* 21(3):329-48.

Selden, Thomas M,. and Michael J. Wasylenko. 1992. “Benefit Incidence Analysis in Developing Countries.” Policy Research Working Papers WPS 1015, Public Economics Department, World Bank.

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

Small, Kenneth, and Harvey Rosen. 1981. “Applied Welfare Economics with Discrete Choice Models.” *Econometrica* 49(1):105-30.

Van de Walle, Dominique. 1998. “Assessing the Welfare Impacts of Public Spending.” *World Development* 26(3):365-79.

Van de Walle, Dominique 1995. “The Distribution of Subsidies through Public Health Services in Indonesia, 1978-87.” In Van de Walle and Nead, op.cit.

Van de Walle, Dominique, and Kimberly Nead. 1995. *Public Spending and the Poor*. Baltimore: The Johns Hopkins University Press.

Yitzhaki, Shlomo, and Slemrod, Joel. 1991. "Welfare Dominance: An Application to Commodity Taxation," *American Economic Review* 81(3):480–96.

Younger, Stephen D. 1999. "The Relative Progressivity of Social Services in Ecuador." *Public Finance Review* 27(3):310-52.

Younger, Stephen D. 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*. Paris: OECD.

TABLE 1. Estimates for the Lanjouw/Ravallion Model Using Province and Individual Data, Cross-section and Panel

		Province-level model, 1994				Household-level model, 1994			
		Coefficient	std error	t-statistic/1	N	Coefficient	std error	t-statistic/1	N
Quintile 1	Intercept	-1.293	1.129	-1.145		9.785	2.249	4.351	
	Slope	0.688	0.136	-2.302	61	0.387	0.059	-10.469	210
Quintile 2	Intercept	0.652	1.400	0.466		3.719	2.893	1.285	
	Slope	0.778	0.179	-1.246	68	0.912	0.079	-1.106	227
Quintile 3	Intercept	-0.219	1.703	-0.128		-9.801	2.861	-3.426	
	Slope	1.241	0.216	1.119	62	1.493	0.084	5.895	258
Quintile 4	Intercept	-0.315	1.606	-0.196		3.136	2.939	1.067	
	Slope	1.210	0.212	0.989	72	1.003	0.089	0.030	280
Quintile 5	Intercept	1.175	1.652	0.711		-6.839	2.825	-2.421	
	Slope	1.084	0.205	0.407	67	1.205	0.091	2.267	363

		Province-level model, 1997				Province-level model, 1994-97 panel/2			
		Coefficient	std error	t-statistic/1	N	Coefficient	std error	t-statistic/1	N
Quintile 1	Intercept	2.950	1.689	1.747					
	Slope	0.250	0.218	-3.443	62	0.611	0.217	-1.797	38
Quintile 2	Intercept	1.067	1.612	0.662					
	Slope	0.888	0.211	-0.532	78	0.595	0.417	-0.972	49
Quintile 3	Intercept	-0.179	1.892	-0.095					
	Slope	1.039	0.249	0.158	76	0.908	0.711	-0.130	44
Quintile 4	Intercept	-2.019	1.751	-1.153					
	Slope	1.454	0.230	1.976	76	0.579	0.479	-0.880	50
Quintile 5	Intercept	-1.820	1.799	-1.011					
	Slope	1.369	0.242	1.528	77	2.307	1.735	0.753	48

/1 The t-statistics test against zero for the intercept and one for the slope.
/2 Intercept coefficients are province-specific in the panel, and thus suppressed.
Estimates are for constrained models, with intercepts summing to zero and slopes to 5.

TABLE 2. Probit Estimates for Secondary School Choice in Rural Peru, 1994

	Regression results			Data	
Household characteristics	coeff	std error	t-statistic	Mean	std error
Constant	-1.4144	0.9083	-1.56	1.00	0.00
Net expenditures/1000[1/]	0.7007	0.2606	2.69	7.81	5.58
Net expenditures/1000 squared[1/]	-0.0037	0.0030	-1.23	92.16	191.82
Net expenditures X distance	0.0016	0.0017	0.98	34.38	53.31
Age of HH head/10	0.1029	0.0463	2.22	4.70	1.18
Gender of HH head	0.0250	0.1400	0.18	0.09	0.29
Head born in urban area	0.2934	0.1019	2.88	0.31	0.46
HH head years of school/10	1.6635	0.3223	5.16	0.43	0.35
HH head years of school/10 squared	-0.4724	0.2249	-2.10	0.31	0.45
HH members age 0 to 5	-0.1262	0.0456	-2.77	0.92	1.03
HH members age 6 to 12	-0.1691	0.0485	-3.49	0.95	0.94
HH members age 13 to 18	0.1259	0.0491	2.57	1.58	0.97
HH members age 19 to 60	0.0704	0.0376	1.87	2.45	1.20
HH members over 60	-0.0465	0.0948	-0.49	0.22	0.51
Child characteristics					
Age/10	-0.0909	0.9591	-0.09	1.58	0.31
Age/10 squared	0.3891	0.2772	1.40	2.58	1.10
Gender	-0.1558	0.0948	-1.64	0.50	0.50
Indigenous	-0.3083	0.1630	-1.89	0.32	0.47
Indigenous X gender	-0.0061	0.1546	-0.04	0.16	0.37
Married	-0.3689	0.3453	-1.07	0.09	0.28
Married X gender	-1.9072	0.5291	-3.60	0.07	0.25
Child of HH head	0.4977	0.4975	1.00	0.83	0.38
Spouse of HH head	1.1789	0.6529	1.81	0.04	0.20
Other HH member	0.4056	0.5061	0.80	0.12	0.32
School characteristics					
Number of required books	0.1287	0.0472	2.73	1.71	0.91
Distance to school	-0.0945	0.0266	-3.56	4.41	7.03
Distance squared	0.0013	0.0005	2.53	68.85	1312.05
Time to secondary school	-0.0020	0.0017	-1.18	37.64	36.69
Primary repetition rate	-0.4951	0.2305	-2.15	0.36	0.20
Secondary repetition rate	1.4337	0.6125	2.34	0.03	0.06
Share of parents expressing a desire to improve:					
School building	-0.0445	0.2126	-0.21	0.66	0.21
Desks and services	-0.3150	0.1324	-2.38	0.48	0.32
Feeding programs	-0.1293	0.2100	-0.62	0.23	0.21
Class size	-0.2476	0.3831	-0.65	0.05	0.09
Teacher training	-0.2070	0.1708	-1.21	0.38	0.26
Teaching materials	-0.2743	0.2549	-1.08	0.17	0.15
Library	-0.2363	0.1856	-1.27	0.32	0.23
Director's power	0.1301	0.6849	0.19	0.02	0.06
Auxiliary personnel training	-0.3209	0.3508	-0.91	0.04	0.08
Other	0.0759	0.2022	0.38	0.26	0.24

[1/] The net expenditure coefficients are constrained to be equal across the school/no school options. All others are for the differential utility of choosing the schooling option.

TABLE 3. Estimated quintile shares of marginal benefits to secondary schooling in rural Peru, 1994

Method:[2]	1a	1b	1c	2	3a[3]	3a[4]	3b[3]	3b[4]	4
	(A)	(B)	(C)	(D)	(E)	(F)	(G)	(H)	(I)
Quintile									
1	0.14	0.12	0.08	0.36	0.16	0.13	0.21	0.15	0.10
	(0.027)	(0.043)	(0.012)	(1.035)	(0.048)	(0.010)	(0.023)	(0.010)	(0.014)
2	0.16	0.12	0.18	0.81	0.16	0.18	0.22	0.20	0.17
	(0.036)	(0.083)	(0.016)	(1.824)	(0.047)	(0.011)	(0.026)	(0.011)	(0.016)
3	0.25	0.18	0.30	-0.19	0.20	0.21	0.22	0.22	0.23
	(0.043)	(0.142)	(0.017)	(1.425)	(0.051)	(0.012)	(0.026)	(0.012)	(0.019)
4	0.24	0.12	0.20	0.44	0.20	0.24	0.20	0.24	0.24
	(0.042)	(0.096)	(0.018)	(1.149)	(0.056)	(0.013)	(0.027)	(0.012)	(0.020)
5	0.22	0.46	0.24	-0.43	0.29	0.24	0.15	0.20	0.26
	(0.041)	(0.347)	(0.018)	(1.880)	(0.055)	(0.015)	(0.022)	(0.012)	(0.022)

Notes: [1] Numbers in parentheses are standard errors.

[2] Methods are as follows:

1a: Lanjouw-Ravallion (1999) applied to 1994 cross-section of province-level data

1b: same as 1a, but applied to a panel of provinces, 1994 and 1997, with fixed effects.

1c: same as 1a, but applied to individual level data.

2: shares of observed changes in secondary school attendance, 1994 to 1997.

3a: shares of estimated compensating variations associated with a policy change, 1994 individual-level data.

3b: shares of estimated change in probability of attendance associated with a policy change, 1994 individual-level data.

4: shares based on standard benefit incidence, using a 0/1 indicator of attendance.

[3] Simulation of a reduction in distance to school to a maximum of two kilometers.

[4] Simulation of a reduction in school fees of 100 soles