

Benefits on the Margin: Observations on Marginal Benefit Incidence

Stephen D. Younger

Benefit incidence analysis has become a popular tool over the past decade, especially for researchers at the World Bank. Despite or perhaps because of the popularity of this method, recent research has pointed out many of its limitations. One of the most common criticisms of benefit incidence analysis is that its description of average participation rates is not necessarily useful in guiding marginal changes in public spending policies. This article considers a variety of methods for analyzing the marginal benefit incidence of policy changes. A key conceptual point is that despite the fact that the various methods measure “marginal” incidence, they do not 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 of interest for some analyses and inappropriate for others. Empirically, the precision of the methods differs substantially, with those relying on differenced data or aggregations of households yielding standard errors that are quite large relative to the estimated shares.

The past decade has seen a resurgence of interest in the relationship between poverty and public spending in developing economies. This resurgence has fostered the return of incidence analysis, particularly for the benefits of public spending in the social sectors. Although analysis of tax incidence has a long and venerable history in economics, distributional analysis of the benefits of public spending—and public policy more generally—is more recent (Aaron and McGuire 1970; Brennan 1976; Meerman 1979; Selowsky 1979). Broadly stated, benefit incidence analysis assesses how the benefits of government spending are distributed across the population.

Though there are many ways to approach this issue, a fairly standard method has emerged, largely based on the work of researchers at the World Bank (Demery 1997; van de Walle and Nead 1995; Selden and Wasylenko 1992). This method takes “across the population” to mean “across the expenditure (or income) dis-

Stephen D. Younger is Associate Director of the Food and Nutrition Policy Program at Cornell University. His e-mail address is sdyl@cornell.edu. The author is grateful for helpful comments from François Bourguignon, Peter Glick, and participants in the Organisation for Economic Co-operation and Development research program “Development of Human Resources and Poverty Alleviation.” Some of the research reported in this article was supported by the U.S. Agency for International Development under cooperative agreement AOT-0546-A-00-3178-00.

DOI: 10.1093/wber/lhg009

© 2003 The International Bank for Reconstruction and Development / THE WORLD BANK

tribution”—an approach consistent with the overall concern with poverty. It then estimates the distribution of benefits based on some variant of the average participation rate in a public program among people in different expenditure brackets.

Given a presumed preference for public spending that benefits the poor, programs and policies are viewed more favorably if poor people's average participation in them is higher than that of the nonpoor. A large increase in the availability of nationally representative, multipurpose surveys—such as the Living Standards Measurement Surveys (Grosh and Glewwe 1998)—and the relative ease with which this standard method can be applied have led to a profusion of such analysis. Indeed, benefit incidence analysis has become a common feature of developing countries' poverty profiles and of many project proposals and evaluations.

Despite or perhaps because of the popularity of this method, recent research has pointed out many of its limitations (van de Walle 1998; Lanjouw and Ravallion 1999). Among the most common criticisms of standard benefit incidence analysis is that its description of average participation rates is not necessarily useful in guiding marginal changes in public spending policies—a point first made by Lipton and Ravallion (1995). The logic of this argument is compelling. The standard method describes who is currently benefiting from a particular public expenditure. As such, it is a useful guide to the effects of a policy change that distributes benefits in proportion to current benefits.

But a policy change that increases spending will not necessarily go to existing beneficiaries in proportion to their current benefits—or even go to existing beneficiaries at all. Many policies explicitly aim to expand the benefits of public spending among nonbeneficiaries. In such cases, because the benefits do not go to existing beneficiaries, the standard method is misleading. Even if services change for existing beneficiaries, the changes may not be uniform, in which case the standard method is also inappropriate. For example, a policy to ensure that all students have a complete set of textbooks will have different distributional consequences if some students already have a complete set, and so gain nothing, whereas others do not.

In response to such observations, several recent studies have proposed alternative methods to measure the marginal incidence of public spending. Glick and Razakamanantsoa (2001) and Younger (2002) examine shares of the change in benefits over time across the expenditure distribution. Lanjouw and Ravallion (1999) and Galasso and Ravallion (2001) estimate the “marginal odds of participation” for each expenditure quantile as the coefficient in a regression of quantile and small area participation rates on large area participation rates. Lanjouw and others (2002) and Ravallion (1999) apply similar techniques to panel data 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 article considers each of these approaches to analyzing the marginal benefit incidence of policy changes using a specific example of secondary educa-

tion in rural Peru. A key conceptual point is that despite the fact that all of these methods claim to measure “marginal” incidence, they do not 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 appropriate for some analyses and inappropriate for others. Empirically, the precision of the methods differs substantially, with those that rely 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 those in the Peru surveys, which are about the same size as many existing multipurpose household surveys.

I. METHODS

This section presents six alternative methods to measure the marginal incidence of a public policy change and compares them to the standard benefit incidence method.

The Standard Benefit Incidence Method

A standard benefit incidence study requires two components: a measure of the value of the benefit that an individual, household, or population group receives from a particular public expenditure; and a way to compare the beneficiaries to the general population. When studying the benefits of public services, the standard method usually uses the government’s cost of provision to estimate a 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—that is, a user or beneficiary gets a benefit of one, others get zero.

It is possible to compare the beneficiaries of a public expenditure with the general population along many dimensions, including ethnicity, gender, region of residence, age, functional income classifications, or political constituency. But an interest in poverty and inequality implies that most comparisons will involve welfare. That is, the goal is to know how the recipients’ welfare compares with the general population’s welfare. Almost all work analyzing developing economies uses household expenditures per capita or per adult equivalent as its measure of welfare. Once it has been 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 or her share of benefits, and sum those shares across a population group, usually welfare quantiles.

The standard method clearly uses group averages to estimate the distribution of benefits. Nevertheless, this average measure does yield the distributional consequences of a marginal policy change that distributes benefits to existing users in proportion to their benefit. An obvious example of such a policy change is a tax or subsidy that changes an existing price proportionately. But one can think of others, such as a new uniform for each child in school or a new vaccination

for 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, policymakers often do not think of increases in public spending 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 or reduced rationing rather than reduced price. As noted, by definition such benefits do not go to existing beneficiaries, so the standard method is inappropriate. The next section discusses several methods for estimating the distributional effects of nonproportional expansions of public service coverage.

Estimating the Benefits of a Marginal Expansion in Services

METHOD 1A: USING SPATIAL VARIATIONS IN COVERAGE TO ESTIMATE MARGINAL PROGRAM BENEFITS. Lanjouw and Ravallion (1999) develop a political economy model in which different population groups—such as the poor and nonpoor—have different political power and different costs and benefits from a given public expenditure. (Similar models are found in Ravallion 1999, 2002.) The interplay between these factors determines the relationship between the size of a program or service, total public spending on it, and each group’s share of its benefits. “Early capture” by the poor occurs when they receive larger shares of a small program but their share declines as the program grows.¹ “Late capture” is the opposite case.

Even with substantial restrictions, the theoretical model yields no general results on whether early or late capture will occur, so the question requires empirical analysis. To that end, Lanjouw and Ravallion estimate the following regression:

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

where i 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 the program participation rate for the department on the participation rates of people in a

1. In Lanjouw and Ravallion’s specification, the nonpoor bear all the program costs and hold all the political power in the sense that the poor cannot impose on them a program that lowers their welfare. In such cases the convexity of the program cost function is sufficient to guarantee “early capture” by the poor.

given province and quantile. The regression is run separately for each quantile. In addition, because $p_{i,k,q}$ is included in p_k , there is an upward bias in the estimate of β_q . Lanjouw and Ravallion resolve this by instrumenting p_k with the left-out mean—that is, the participation rate for all of department k except those individuals in province i and quantile q .

The intuition behind the regression is that by observing variations in departmental participation across the country, it is possible to understand how increased coverage affects the participation of different population 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 and quantile. One advantage of this method is that it requires only a cross-section of data, just like the standard method. An important assumption is that across regions, the same political process determines the correlation between program size or coverage and incidence.

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

METHOD 1B: CONTROLLING FOR FIXED EFFECTS. Lanjouw and Ravallion (1999) point out that equation 1 includes no controls for any effects on province and quantile participation rates except the department's participation rate. Where data 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. (Lanjouw and others 2002 and Ravallion 1999 pursue this strategy.) This is possible even if the surveys are not panels of households, as long as the households are sampled from the same provinces and each survey is representative at the province level.²

METHOD 1C: USING DISAGGREGATED (INDIVIDUAL) DATA. A purely statistical problem with the Lanjouw-Ravallion model is that it uses average data for provinces and quantiles. Although this approach was often necessary in the past, individual-level data sets are now available to estimate benefit incidence. Grouping observations into province and quantile averages reduces the efficiency of the estimates, yielding larger estimated standard errors (Johnston 1972). The

2. The first condition is often true of household surveys, whereas the second is quite rare. In the two Peru surveys used in the next section, there is considerable overlap between provinces, 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. In Peru, however, survey teams usually return to the same clusters when they conduct new surveys, which should minimize this problem.

application below estimates the model for both group average and individual data.

METHOD 2: OBSERVING CHANGES AS PROGRAMS EXPAND OVER TIME. This method addresses the same margin as method 1, the incidence of increased spending 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 substitutes the change in a given quantile's benefits (or program participation) for its level. Glick and Razakamanantsoa (2001) and Younger (2002) use two cross-sections at different points in time to estimate each quantile's share in the change in use of various public services. Van de Walle 1995, Hammer and others 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.

This method requires at least two cross-sectional surveys. But just as the number of developing economies with at least one nationally representative multipurpose survey grew over the past decade, an increasing number of countries now have more than one such survey a few years apart. Like method 1, this approach 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.

METHOD 3A: ECONOMETRIC ESTIMATES OF COMPENSATING VARIATIONS. Rather than use the standard benefit incidence method as an approximation for the compensating variation for a price change, it is possible to estimate, econometrically, compensating variations for price and other policy changes. A well-established literature in transport and environmental economics does this for goods and services where demand is discrete (Small and Rosen 1981; McFadden 1995), and Paul Gertler and his associates have applied these techniques to the demand for health and education in developing countries (Gertler and others 1987; Gertler and Glewwe 1990).

The model is well known. Using education as an example, assume that each household has a utility function that depends on its consumption and its choice for the quality of schooling:

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

where j indexes the choice (school or no school), y is household permanent income proxied by household expenditures, p_j is the price of choice j including all opportunity costs of time, and Q is a function that measures quality, which depends on choice-specific characteristics X_j and household or personal characteristics Z . Households choose the option j that yields the highest utility. Although V_j is not observable, if a household chooses option j , V_j is greater than all other

V_j . The model estimates the probability that this is the case using only the observed choice, and takes the probability of choosing option j as an expected demand for that option. Small and Rosen (1981) 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 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 . But that may not be the case for some public services. For example, schools may exclude students based on merit, gender, social status or connections, and so on. In such cases some school-age children may not be attending school even though that option would provide them with the greatest utility, which biases the estimates in equation 2 if the rationing is correlated with the regressors.³ In the case study that follows, rationing is not a problem. Except for 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 policy variable x_j is changed?" 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 the 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.

II. BENEFIT INCIDENCE OF SECONDARY SCHOOLING IN RURAL PERU

This section estimates the distribution of benefits from an expansion of secondary schooling in rural Peru using all the methods outlined in the previous section. In addition, it calculates standard measure of benefit incidence. The calculations for shares in observed changes in participation—that is, method 2—are straightforward. Note that secondary schooling coverage rates increased slightly in rural

3. Some nonprice 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.

Peru between 1994 and 1997, from 6.33 to 6.58 percent of the rural population. The other methods require preliminary regressions, which are reported first.

Estimating the Marginal Odds of Participation

The results of four regressions based on the Lanjouw and Ravallion (1999) model, as described in method 1, are reported in table 1 and table 2. The cross-sectional data are for rural households from the 1994 and 1997 rounds of the Encuesta Nacional de Hogares sobre Medición de Niveles de Vida (the National Household Living Standards Measurement Survey). The regressions are for rural provinces only, to be consistent with the model for method 3. In the Individual-level model section of table 1 the dependent variable is a zero/one indicator of individual-level participation in secondary schooling. (Thus, the linear probability model is being estimated; it would also be possible to estimate this model as a probit or logit.) In the other sections of tables 1 and 2 the data for the dependent variable are the province- and quintile-specific participation rates for secondary school attendance, defined as the number of secondary students divided by the population.

In all the 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. In addition, two restrictions have been imposed on the coefficients: that the α_q s sum to zero and the β_q s sum to the number of quantiles, five in this case. Although Lanjouw and Ravallion do not impose these restrictions, they are required if the estimated shares of marginal benefits are to sum to one. As it turns out the unrestricted estimates are quite close to those reported here for all the cross-sectional models. For the panel data model the differences are much larger—but so are the standard errors, so even

TABLE 1. Estimates for the Lanjouw-Ravallion (1999) Model Using Province- and Individual-Level Cross-Sectional Data, 1994

Quintile ^a		Province-level model, 1994				Individual-level model, 1994			
		Coefficient	SE	<i>t</i> -statistic ^b	N	Coefficient	SE	<i>t</i> -statistic ^b	N
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
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
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
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
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

^aQuintiles are based on household expenditures per capita.

^bThe *t*-statistics test against zero for the intercept and one for the slope.

Source: Author's calculations

TABLE 2. Estimates for the Lanjouw-Ravallion (1999) Model Using Province-Level Cross-Sectional and Panel Data, 1994–97

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

^aIntercept coefficients are province-specific in the panel and thus suppressed. Estimates are for constrained models, with intercepts summing to zero and slopes to five.

^bQuintiles are based on household expenditures per capita.

^cThe *t*-statistics test against zero for the intercept and one for the slope.

Source: Author's calculations

for this model the slope coefficients in the restricted and unrestricted models do not differ at the 5 percent level.

Three points about table 1 and table 2 are worth noting here (see later discussion of the distributional implications). First, the various models produce quite different estimates for the quintile-specific marginal odds of participation. This is true even for the two cross-sectional models (Province-level model, 1994 and 1997). These differences cannot be due to changing marginal odds of participation as coverage expands, because in these samples rural secondary school enrollments were essentially constant between 1994 and 1997. Although 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 in the individual-level model are only about half those of the province-level model.

Estimating Demand

Estimates of the demand for secondary schooling in rural Peru, using the same 1994 data set, are shown in table 3. (The 1994 data are used because that survey asked a broader range of questions about school quality, including distance to school and parents' evaluations of problems at their children's schools.) Though it is customary to consider multiple options for schooling—no school, public school, or private school, or no school, local school, or distant school—in rural Peru only 3 percent of children attend private school, making an estimate of the demand for private options infeasible. Moreover, the survey does not identify students who are at school away from home. Thus, the table esti-

TABLE 3. Probit Estimates for Secondary School Choice in Rural Peru, 1994

Variable	Regression results			Data	
	Coefficient	SE	t-statistic	Mean	SE
<i>Household characteristics</i>					
Constant	-1.4144	0.9083	-1.56	1.00	0.00
Net expenditures/1000 ^a	0.7007	0.2606	2.69	7.81	5.58
Net expenditures/1000 squared ^a	-0.0037	0.0030	-1.23	92.16	191.82
Net expenditures × distance	0.0016	0.0017	0.98	34.38	53.31
Age of household head/10	0.1029	0.0463	2.22	4.70	1.18
Gender of household head	0.0250	0.1400	0.18	0.09	0.29
Household head born in urban area	0.2934	0.1019	2.88	0.31	0.46
Household head years of schooling/10	1.6635	0.3223	5.16	0.43	0.35
Household head years of schooling/10 squared	-0.4724	0.2249	-2.10	0.31	0.45
Household members ages 0–5	-0.1262	0.0456	-2.77	0.92	1.03
Household members ages 6–12	-0.1691	0.0485	-3.49	0.95	0.94
Household members ages 13–18	0.1259	0.0491	2.57	1.58	0.97
Household members ages 19–60	0.0704	0.0376	1.87	2.45	1.20
Household members over 60	-0.0465	0.0948	-0.49	0.22	0.51
<i>Child characteristics</i>					
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 × gender	-0.0061	0.1546	-0.04	0.16	0.37
Married	-0.3689	0.3453	-1.07	0.09	0.28
Married × gender	-1.9072	0.5291	-3.60	0.07	0.25
Child of household head	0.4977	0.4975	1.00	0.83	0.38
Spouse of household head	1.1789	0.6529	1.81	0.04	0.20
Other household member	0.4056	0.5061	0.80	0.12	0.32
<i>School characteristics</i>					
Number of required textbooks	0.1287	0.0472	2.73	1.71	0.91
Distance to school	-0.0945	0.0266	-3.56	4.41	7.03
Distance to school squared	0.0013	0.0005	2.53	68.85	1312.05
Time to 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
<i>Percentage of parents expressing a desire to improve</i>					
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

^aThese coefficients are constrained to be equal across the schooling or no schooling options. All others are for the differential utility of choosing the schooling option.

Source: Author's calculations

mates only the choice between attending or not attending school, as a probit. The analysis is limited to rural areas because a model with only a few choices is inappropriate for most urban areas. A resident of Lima (Peru's capital) has a choice of many schools, public or private. No survey in Peru makes it possible to adequately identify (let alone model) these choices.

Because the probit can identify the model only up to the differences in V_j , the model must normalize against one option, which here will be the no school choice. Thus, it is assumed that $Q(X_0, Z) = 0$. In the estimates it is assumed that the function $f()$ in equation 2 is quadratic in net expenditures and constrained to be the same for each option. There is some debate about the second 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 from school. This variable is the one that will be used to compare results with the Lanjouw-Ravallion method, so the aim is to allow as much flexibility as possible.⁴

The sample includes all rural children attending secondary school or eligible to attend. The latter group includes all children of the appropriate age who have not graduated from secondary school—including children who have not graduated from primary school, because in the context of long-run optimization the decision not to complete primary school is partly affected by perceptions of the value of secondary school. Dow (1999) defends this type of unconditional estimate.

All the household characteristics in table 3 are self-explanatory except for net household expenditures. Household expenditures are defined in the broadest way possible, including imputed values of owner-occupied housing and own-produced goods (Younger 2002). In addition, if a household contains a secondary student, the costs of schooling—including fees, books, uniforms, and transportation—and the opportunity costs of time at school are added to get a gross expenditure variable that is before the costs of schooling. The price is the cluster- or district-level mean for these school costs. (Districts are the third-level geopolitical unit in Peru, smaller than provinces.) If a cluster has at least four observations, cluster-level data are used to calculate means. Otherwise district-level data are used. (The same criterion of at least four observations is used for all subsequent cluster- or district-level regressors.) All means calculated in this way are left-out means. Net expenditures are gross expenditures minus this price variable.

The child characteristics in table 3 are also self-explanatory. The default option for relation to the household head is being the head. The school characteristics require some explanation. The number of required books is the cluster- or

4. 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.

district-level mean number of textbooks that school requires. Here this is taken as an indicator of academic quality. Distance to the school is measured 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 a list of school characteristics. The regressors in table 3 are the cluster- or district-level share of parents expressing in interest in improving each feature.

The signs of the coefficients are almost all consistent with prior expectations. Net household expenditures have a positive and only slightly concave effect on the probability of secondary school attendance in rural areas. Children are more likely to attend secondary school if they live in a household whose head is older, urban-born, and more educated. Children in households with younger children (age 12 and under) are less likely to attend secondary school, whereas those in households with older children (excluding the child being observed) or adults are more likely to attend.

Of the child characteristics, only one variable is statistically significant at standard levels: married girls are less likely to attend secondary school even after accounting for the positive effect resulting 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 attendance. On the other hand, only one of the parental opinion questions is significantly different from zero—that expressing a desire for better desks and services.

The results in table 3 are used to simulate two policy changes. The first is a reduction in school 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 second 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 one 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 most consistent with methods 1 and 2. For each simulation calculations are made for each child of the compensating variations for the policy change (method 3A) and of the change in the probability that he or she attends school (method 3B).

McFadden (1995) shows that the standard method for calculating compensating variations in the discrete choice model developed by Small and Rosen (1981) will yield biased estimates if utility is a nonlinear function of income, as it is in this model. So, here the compensating variations are calculated using the simulation method described by McFadden with 1000 repetitions for each observation.

Results

The estimated marginal benefits from secondary schooling to each per capita expenditure quintile in rural Peru, calculated using each of the methods presented—along with method 4, which is a standard benefit incidence analysis—are shown in table 4.⁵ The shares for the Lanjouw and Ravallion (1999) methods are those presented in tables 1 and 2 divided by five.

Consider first the estimated distributions of marginal benefits associated with a program expansion—that is, methods 1A, 1B, 1C, and 2 in columns A through D. As tables 1 and 2 showed, the results for the Lanjouw-Ravallion methods (1A, 1B, 1C) are statistically different from equal shares, with a modest antipoor bias. (All tests are at the 5 percent confidence level.) Only a few of the estimates differ by economically important amounts, but except for the difference in the first quintile for methods 1A and 1C, these differences are not statistically significant due to the relatively large standard errors. This is especially the case for the panel data method (1B), which estimates a much larger marginal share for the fifth 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 in 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 is possible because a quintile can have negative marginal benefits—that is, a decline in its participation rate over time—even as overall participation increases. (This was true for the third and fifth quintiles between 1994 and 1997 in rural Peru.) Method 2 seems unsatisfactory on two counts. First, the 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 secondary school attendance 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, method 2 is meant to capture shares of marginal changes, which by definition are small. Thus, applying this method to services with larger expansions might produce more stable estimates, but such estimates 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. Though this is only one example, it seems that

5. The quintiles are based only on the rural samples of the Peru surveys. It would be just as easy to derive them for all households in the sample and give zero benefits to urban residents. Where rural residents fall in the nationwide spending, 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.

TABLE 4. Quintile Shares of Marginal Benefits to Secondary Schooling in Rural Peru, Calculated Using Various Methods

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

Note: Numbers in parentheses are standard errors. The methods are 1A: Lanjouw and Ravallion (1999) model applied to 1994 cross-section of province-level data. 1B: same as 1A but applied to a panel of provinces, in 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, 1994-97. 3A: shares of estimated compensating variations associated with a policy change, individual-level data, 1994. 3B: shares of estimated change in probability of attendance associated with a policy change, individual-level data, 1994. 4: shares based on standard benefit incidence, using a zero or one indicator of attendance. Quintiles are based on household expenditures per capita.

^aSimulation of a reduction in distance to school to a maximum of 2 km.

^bSimulation of a reduction in school fees of 100 soles.

Source: Author's calculations.

method 2 may require samples that are much larger than are typical in developing economies to produce precise estimates of marginal shares.

As noted, increased access is the policy that most people have in mind when considering a program expansion. For secondary schools in rural Peru, this is best captured by reduced distance to school—a relevant variable in Peru during the period being analyzed. In the 1990s Peru's government invested considerable resources in building and rehabilitating schools through the Fondo Nacional de Compensación y Desarrollo Social (the National Compensation and Social Development Fund) (Paxson and Schady 1999), and in rural areas the median travel time to get to school fell from 40 to 30 min between 1994 and 1997. (The 1997 survey did not ask for distance to school, only travel time.) 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 during this period. As such, it is interesting to compare method 1 to method 3.

As noted, the policy change simulated here is a reduction in distance to school to a maximum of 2 km.⁶ The estimated distributions of the changes in the probability of attendance (column G, representing method 3B, in table 4) are somewhat more progressive than those for the compensating variations (column E, representing method 3A), though only the fifth quintile's shares differ significantly. These findings indicate that in rural Peru, the value that people place on secondary schooling increases with household expenditures per capita. Both estimates of marginal shares are somewhat larger than any of the Lanjouw-Ravallion methods for the poorest quintile.

For the changes in the probability of attendance (column G) all of these differences are statistically significant. But for the compensating variations (column E) only the first quintile difference with method 1C—the most precise of the Lanjouw-Ravallion methods—is statistically significant. The changes in probability (column G) also differ significantly from method 1C (the individual-model) in the third and fifth quintiles, and from method 1A (the province-level cross-section) in the fifth quintile. For the compensating variations (column E) the only other significant difference is for the shares for the third quintile relative to method 1C. 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 in table 4) based on school attendance are quite close to those derived from both the change in probability of attendance (column H, representing method 3B) and the compensating variations (column F, representing method 3A) for a 100 sole price change—though the differences at the first quintile are statistically significant, if minor (0.05 and 0.03, respectively), as is the differ-

6. The mean distance in the sample is 4.4 km and the median is 2.7 km, with a standard deviation of 7.0 km. The extreme distance is 36.0 km.

ence at the fifth quintile for model 3B (column H). 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 (1999) result—the shares from the standard benefit incidence method for rural secondary school attendance in Peru are not significantly less propoor than those from any of the other methods, except for the simulation of reduced distance to school.

III. CONCLUSION

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, whereas analysis of any policy change should be done on the margin. This article argues that the standard method can in fact be interpreted as a marginal method: it gives a first-order approximation of the distributional consequences of a price change or any other change that affects only observed beneficiaries in proportion to their existing benefit. In the example of secondary school attendance in rural Peru, the approximation is reasonably good, even for a large (nonmarginal) change in the cost of attendance. This finding is consistent with previous work on five social services in Ecuador (Younger 1999). In that sense this article supports the standard method—as long as it is interpreted correctly.

More broadly, however, there are more margins of interest than price. In particular, expanded access to services, rather than changes in fees, is often what policymakers have in mind when considering increased spending on a public service. This article 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 a household's compensating variations or willingness to pay for that change—or, more narrowly, on the change in its probability of participation.

In general, the different methods produce different estimates of marginal benefit incidence, suggesting that analysts should tailor their choice of method to the issue 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 in program coverage in an unchanging political-economic environment, one of the Lanjouw-Ravallion methods (method 1) is the relevant option. Those interested in analyzing specific policy

changes with effects that are 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, methods 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 data is preferable where such data are available. Second, methods that rely on differences over time have large standard errors. Though 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." World Bank, Poverty Reduction and Economic Management Network, Poverty Anchor, Washington, D.C.
- 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." World Bank, Development Research Group, Washington, D.C.
- 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, Luis 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 University, Cornell Food and Nutrition Policy Program, Ithaca, N.Y.
- 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 University, Cornell Food and Nutrition Policy Program, Ithaca, N.Y.
- 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 Dominique Van de Walle and Kimberley Nead, eds., *Public Spending and the Poor*. Baltimore, Md.: Johns Hopkins University Press.
- 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." *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: Organisation for Economic Co-operation and Development.
- Lipton, Michael, and Martin Ravallion. 1995. "Poverty and Policy." In Jere Behrman and T. N. Srinivasan, eds., *The Handbook of Development Economics*, vol. 3B. Amsterdam: North-Holland.
- McFadden, Daniel. 1995. "Computing Willingness-to-Pay in Random Utility Models." University of California at Berkeley, Department of Economics.
- 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)." Policy Research Working Paper 2229. World Bank, Washington, D.C.
- Ravallion, Martin. 1999. "Is More Targeting Consistent with Less Spending?" *International Tax and Public Finance* 6(3):411-19.
- . 2002. "Who Is Protected? On the Incidence of Fiscal Adjustment." Paper presented at the International Monetary Fund Conference on Macroeconomic Policies and Poverty Reduction, Washington, D.C. March.
- 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 Paper 1015. World Bank, Washington, D.C.
- 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. 1995. "The Distribution of Subsidies through Public Health Services in Indonesia, 1978-87." In Dominique Van de Walle and Kimberley Nead, eds., *Public Spending and the Poor*. Baltimore, Md.: Johns Hopkins University Press.
- . 1998. "Assessing the Welfare Impacts of Public Spending." *World Development* 26(3):365-79.
- Van de Walle, Dominique, and Kimberley Nead. 1995. *Public Spending and the Poor*. Baltimore, Md.: Johns Hopkins University Press.
- Yitzhaki, Shlomo, and Joel Slemrod. 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.
- . 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: Organisation for Economic Co-operation and Development.