Robust Multidimensional Poverty Comparisons

by

Jean-Yves Duclos
Département d'Économique and CIRPÉE, Université Laval, Canada,
and Institut d'Anàlisi Econòmica (CSIC), UAB, Barcelona, Spain

David E. Sahn
Cornell University

and

Stephen D. Younger
Cornell University

Abstract

We demonstrate how to make poverty comparisons using multidimensional indicators of well-being, showing in particular how to check whether the comparisons are robust to aggregation procedures and to the choice of multidimensional poverty lines. In contrast to earlier work, our methodology applies equally well to what can be defined as "union", "intersection" or "intermediate" approaches to dealing with multidimensional indicators of well-being. To make this procedure of some practical usefulness, the paper also derives the sampling distribution of various multidimensional poverty estimators, including estimators of the "critical" poverty frontiers outside which multidimensional poverty comparisons can no longer be deemed ethically robust. The results are illustrated using data from a number of developing countries.

Keywords Multidimensional Poverty, Stochastic Dominance. **JEL #** D31,D63,I31,I32.

Duclos' research was supported by SSHRC, FCAR and IDRC's MIMAP programme. Sahn's and Younger's participation was supported by the USAID and AERC. An early version of this work was written as a supporting paper for the World Development Report (2000), regarding which we are grateful to Christian Grootaert for his comments and support. We also thank Rob Strawderman and Steve Westin for technical advice, Nicolas Beaulieu and Wilson Perez for their excellent research assistance, and Peter Lambert, Paul Makdissi and Lars Osberg for very useful comments.

Corresponding addresses:
Jean-Yves Duclos: Département d'économique, Pavillon de Sève, Université Laval, Québec, Canada, G1K 7P4; Tel.: (418) 656-7096; Email: jduc@ecn.ulaval.ca.
David Sahn and Stephen Younger: Room 3M12 MVR Hall, Cornell University, Ithaca, NY 15843. E-mail: des16@cornell.edu or sdy1@cornell.edu.

April, 2005

1 Introduction

Since the seminal work of Amartya Sen (1979, 1985, 1987), it is common to assert that poverty is a multi-dimensional phenomenon, a proposition that most economists accept in theory. Yet in practice, the vast majority of empirical work on poverty uses a one-dimensional yardstick to judge a person's well-being, usually expenditures or income *per capita* or per adult equivalent. Our purpose in this paper is to show that it is possible to bring the empirical literature closer to the theoretical rhetoric by making quite general poverty comparisons when deprivation is measured in multiple dimensions.

We note at the outset that there is a branch of the poverty measurement literature that considers multiple dimensions of well-being, but these papers invariably aggregate the multilple measures of well-being into a one-dimensional index, essentially returning to a univariate analysis. The best-known example is the Human Development Index (HDI) of the UNDP (1990), which uses a weighted average of life expectancy, literacy, and GDP per capita for a population, though several more have been proposed recently.[1] Any such index requires a specific aggregation rule to sum up the components of the index, and any such rule is necessarily arbitrary. This leaves open the possibility that two equally valid rules for aggregating across several dimensions of well-being could lead to contradictory conclusions about which groups have higher poverty.[2]

To avoid that, we develop poverty comparisons that are valid for a broad class of aggregations rules. This is in the spirit of the dominance approach to poverty comparisons, as initially developed by Atkinson (1987) and Foster and Shorrocks (1988a,b,c) in a uni-dimensional context [3]. It is well-known that one important advantage of this approach is that it is capable of generating poverty orderings that are robust to the choice of a poverty index over broad classes of indices – the orderings are "poverty-measure robust". Such comparisons thus relieve the analyst or policymaker from the task of choosing one particular (and ethically arbitrary) poverty measure. To understand our generalization of

[1] See, for example the papers presented at a recent WIDER conference on Inequality, Poverty, and Human Well-Being, http://www.wider.unu.edu/conference/conference-2003-2/conference2003-2.htm.

[2] Such rules have been the focus of some of the recent literature: see for instance Tsui (2002), Dutta, Pattanaik, and Xu (2003), and Bourguignon and Chakravarty (2003). Bourguignon and Chakravarty (2003) also give several interesting examples in which poverty orderings vary with the choice of aggregation rules.

[3] Atkinson and Bourguignon (1982,1987) and Bourguignon (1989) first used this approach in the context of multidimensional social welfare. See also Crawford (1999) for a recent application.

dominance comparisons to multiple dimensions of well-being, it is helpful to think of a one-dimensional poverty index as an aggregation rule within that dimension. That is, the index aggregates all of the information in an income distribution into a single scalar poverty measure. Univariate dominance comparisons are valid for broad classes of such rules, *i.e.*, broad classes of poverty measures. Our multidimensional comparisons will be similarly valid for broad classes of aggregation rules used in any one dimension, and also for broad classes of aggregation rules used across dimensions of well-being.

In contrast to earlier work on multidimensional comparisons, our orderings are also ”poverty-line robust”, in the sense of being valid for the choice of *any* poverty frontier over broad ranges. Given the well-known sensitivity of many poverty comparisons to the choice of poverty lines, and the difficulty of choosing the ”right” poverty line, especially for many non-income dimensions of well-being, this would appear to be an important contribution.

For most of the paper, we limit ourselves to the case of two measures of well-being, though we do provide an example of a three-dimensional comparison. In theory, extending our results to more than two dimensions is straightforward. In practice, though, most existing datasets in developing countries are probably not large enough to support tests on more than a few dimensions of well-being. This is because the curse of dimensionality (Bellman, 1961) affects our non-parametric estimators.

We begin in Section 2 with the theory of making multidimensional poverty comparisons. One of the first conceptual challenges of poverty analysis in multiple dimensions is deciding who is ”poor.” We consider this question in Section 2.1. The literature distinguishes between *intersection* and *union* definitions of poverty. If we measure well-being in the dimensions of income and height, say, then a person could be considered poor if her income falls below an income poverty line *or* if her height falls below a height poverty line. We may define this as a *union* definition of multidimensional poverty. An *intersection* definition, however, would consider a person to be poor only if she falls below *both* poverty lines [4]. In contrast to earlier work, the tests that we develop are valid for both definitions – and also valid, in fact, for any choice of intermediate definitions for which the poverty line in one dimension is a function of well-being measured in the other. We also

[4]For further recent discussion of this, see *inter alia* Bourguignon and Chakravarty (2002,2003), Atkinson (2002) and Tsui(2002).

show how the concept of a poverty line in one dimension extends naturally to a "poverty frontier" in multiple dimensions, with the shape of that frontier determined by whether we are interested in union, intersection, or intermediate poverty measures.

Section 2.2 presents our main theorem for multidimensional poverty comparisons that are first-order in each dimension[5]. Section 2.3 extends these results to higher-order poverty comparisons.

Section 2.5 presents a different approach to multidimensional poverty comparisons. Rather than asking, "Is poverty lower for A than for B over all reasonable poverty frontiers?" we ask, "What is the area of poverty frontiers over which we can be sure that poverty is lower for A than for B?" This approach provides one useful way to get around the need to make an arbitrary choice of "reasonable" limits for the range of admissible poverty frontiers. The procedure also makes the dominance conditions "locate the disagreements that are crucial" (Atkinson (2002)).

While much of the paper is a contribution to the theory of multidimensional poverty measurement, the results are also important for applied poverty analysts. The Millenium Development Goals, for example, focus attention on deprivation in multiple dimensions. Many policy makers take these goals seriously, and they have clearly helped to broaden the development policy debate beyond a narrow focus on reducing income poverty. The methods that we propose here should provide applied researchers with an attractive tool for analyzing these broader definitions of poverty. To demonstrate that, Section 4 gives a series of examples that highlight some of the subtleties of the paper's theoretical results, using data from several developing countries.

One of the key points in this section is that it is possible for a set of univariate analyses done independently for each dimension of well-being to conclude that poverty in A is lower than poverty in B while a multivariate analysis concludes the opposite, and vice-versa [6]. The key to these possibilities is the interaction of the various dimensions of well-being in the poverty measure, and their correlation in the sampled populations. We argue that a reasonable poverty measure should allow the level of deprivation in one

[5] We will define this precisely in Section 2.2, but it is analogous to first-order poverty comparisons in one dimension.

[6] To our knowledge, the literature does not include poverty analyses that look at several dimensions of well-being "one-at-a-time," but the rise of the MDGs could well bring that about as analysts first study income poverty, then health poverty, then education poverty, etc.

dimension to affect our assessment of how much poverty declines if there is an improvement in another dimension. For example, we might think that an increase in income for a poor person should cause a larger decline in a poverty measure if the recipient is also relatively deprived in the dimensions of health, education, *etc*. "One-at-a-time" comparisons of poverty in the dimensions of income, education, and health cannot capture these interdependencies, while our multidimensional measures can. In practice, populations exhibiting higher correlations between measures of well-being will be poorer than those that do not, relative to what one would expect from making univariate comparisons alone.

A final output of the paper is to provide the sampling distribution of many of the estimators that are useful for multidimensional poverty analysis, in such a way that one may infer from sample estimates the true population value of poverty measures. Previous work on multidimensional poverty comparisons has ignored sampling variability, yet this is fundamental if the study of multidimensional poverty comparisons is to have any practical application.

2 Multiple indicators of well-being

2.1 Poverty frontiers and poverty aggregation in two dimensions

Let x and y be two indicators of individual well-being[7]. These could be, for instance, income, expenditures, caloric consumption, life expectancy, height, cognitive ability, the extent of personal safety and freedom, *etc.* Denote by

$$\lambda(x,y) : \Re^2 \rightarrow \Re \left| \frac{\partial \lambda(x,y)}{\partial x} \geq 0, \frac{\partial \lambda(x,y)}{\partial y} \geq 0 \right. \tag{1}$$

a summary indicator of individual well-being, analogous to but not necessarily the same as a utility function. Note that the derivative conditions in (1) simply mean that different indicators can each contribute to overall well-being. We make the differentiability assumptions for expositional simplicity, but they are not strictly necessary so long as $\lambda(x,y)$ is non-decreasing over x and y. It should be clear that the weak inequalities on these derivatives impose few constraints on the precise value of these contributions.

[7]For expositional simplicity, we focus on the case of two dimensions of individual well-being. We will illustrate in section 4 the extension to more than two dimensions.

We then assume that an unknown poverty frontier separates the poor from the rich. We can think of this frontier as a series of points at which overall well-being of an individual is precisely equal to a "poverty level" of well-being, and below which individuals are in poverty. This frontier is assumed to be defined implicitly by a locus of the form $\lambda(x, y) = 0$, and is analogous to the usual downward-sloping indifference curves on the (x, y) space. The set of the poor is then obtained as:

$$\Lambda(\lambda) = \{(x, y) \,|(\lambda(x, y) \leq 0\} . \tag{2}$$

Consider Figure 1 with poverty lines z_x and z_y in dimensions of well-being x and y. $\lambda_1(x, y)$ gives an "intersection" poverty index: it considers someone to be in poverty only if she is poor in *both* of the two dimensions of x and y, and therefore if she lies within the dashed rectangle of Figure 1. $\lambda_2(x, y)$ (the L-shaped, dotted line) gives a union poverty index: it considers someone to be in poverty if she is poor in *either* of the two dimensions, and therefore if she lies below or to the left of the dotted line. Finally, $\lambda_3(x, y)$ provides an intermediate approach. Someone can be poor even if $y > z_y$, if her x value is sufficiently low to lie to the left of $\lambda_3(x, y) = 0$. Alternatively, someone can be non-poor even if $y < z_y$ if her x value is sufficiently high to lie to the right of $\lambda_3(x, y) = 0$.[8]

To define multidimensional poverty indices precisely, let the joint distribution of x and y be denoted by $F(x, y)$. For analytical simplicity, we focus in this paper on classes of additive multidimensional poverty indices[9]. An additive poverty index that combines the two dimensions of well-being can be defined generally as $P(\lambda)$,

$$P(\lambda) \;=\; \int\int_{\Lambda(\lambda)} \pi(x, y; \lambda)\, dF(x, y), \tag{3}$$

where $\pi(x, y; \lambda)$ is the contribution to poverty of an individual with well-being indicators x and y. By the well-known "poverty focus axiom" (see for instance Foster (1984)), this is such that

$$\pi(x, y; \lambda) \begin{cases} \geq 0 \text{ if } \lambda(x, y) \leq 0 \\ = 0 \text{ otherwise.} \end{cases} \tag{4}$$

[8]As noted in Equation 1, $\lambda(x, y)$ is a very general function. Specific examples of bivariate poverty measures that lead to a variety of such functions, including the three in Figure 1, can be found in Bourguignon and Chakravarty (2003).

[9]This is a usual assumption, often obtained as a product of a subgroup decomposability axiom – see for instance Foster, Greer and Thorbecke (1984) and Bourguignon and Chakravarty (2002).

The multidimensional headcount is obtained when $\pi(x, y; \lambda) = 1$ whenever $\lambda(x, y) \leq 0$. The π function in equation (4) is thus the weight that the poverty measure attaches to *** someone who is "poor," *i.e.* inside the poverty frontier. That weight could be 1 (for a headcount), but it could take on many other values as well, depending on the poverty measure of interest. By the focus axiom, it has to be zero for those outside the poverty frontier.

2.2 Poverty dominance in two dimensions

With $F(x)$ being the distribution function for x, recall first that the usual unidimensional stochastic dominance curve is given (for x) by ***

$$P^{\alpha}(z) = \int_0^z (z - x)^{\alpha} \, dF(x) \tag{5}$$

for an integer $\alpha \geq 0$. (5) is thus an average of poverty gaps $(z - x)$ to the power α, and it is also the popular Foster, Greer and Thorbecke (1984) (FGT) poverty index. A *bi-dimensional stochastic dominance surface* can then be defined as:

$$P^{\alpha_x, \alpha_y}(z_x, z_y) = \int_0^{z_y} \int_0^{z_x} (z_x - x)^{\alpha_x} (z_y - y)^{\alpha_y} \, dF(x, y) \tag{6}$$

for integers $\alpha_x \geq 0$ and $\alpha_y \geq 0$. We generate the dominance surface by varying the poverty lines z_x and z_y over an appropriately chosen domain, with the height of the surface determined by equation 6. This function looks like a two-dimensional generalization of the FGT index, and it could be interpreted as such. It is important to highlight, however, that the poverty comparisons that we make will be valid for broad classes of poverty functions, not this one alone.[10] A second important feature of the dominance surface is *** that it is influenced by the covariance between x and y, the two measures of well-being, because the integrand is multiplicative. Rewriting (6), we find indeed that ***

$$\begin{aligned} P^{\alpha_x, \alpha_y}(z_x, z_y) = & \int_0^{z_y} (z_x - x)^{\alpha_x} \, dF(x) \int_0^{z_x} (z_y - y)^{\alpha_y} \, dF(y) \\ & + \operatorname{cov}\left((z_x - x)^{\alpha_x}, (z_y - y)^{\alpha_y}\right). \end{aligned} \tag{7}$$

The height of the dominance surface is therefore the product of the two unidimensional curves plus the covariance in the poverty gaps in the two dimensions. Thus, the higher the correlation between x and y, the higher the dominance surfaces, other things equal.

[10] One of the more remarkable features of the results that follow is that they are also useful for union poverty definitions, even though the function in (6) is itself an intersection definition.

Our poverty comparisons make use of orders of dominance, s_x in the x and s_y in the y dimensions, which will correspond respectively to $s_x = \alpha_x + 1$ and $s_y = \alpha_y + 1$. The parameters α_x and α_y also capture the aversion to inequality in poverty in the x and in the y dimensions, respectively. $P^{0,0}(z_x, z_y)$ generates a cumulative density function, analogous to a poverty incidence curve in a univariate analysis. $P^{1,0}(z_x, z_y)$ sums the poverty gaps in x (given by $\max(z_x - x, 0)$) for those that are poor in y, and then normalizes this sum by the size of the total population. $P^{1,1}(z_x, z_y)$ sums the product of the poverty gaps in x and in y, again normalized by the size of the total population, and can therefore be thought of as a bidimensional average poverty gap index, with the weights on the poverty gaps in one dimension being the poverty gaps in the other dimension. Analogous interpretations exist for other combinations of α_x and α_y values.

To describe the class of poverty measures for which the dominance surfaces defined in equation 6 are sufficient to establish poverty orderings, assume first that the general poverty index in (3) is left differentiable[11] with respect to x and y over the set $\Lambda(\lambda)$, up to the relevant orders of dominance, s_x for derivatives with respect to x and s_y for derivatives with respect to y. Denote by π^x the first derivative[12] of $\pi(x, y; \lambda)$ with respect to x; by π^y the first derivative of $\pi(x, y; \lambda)$ with respect to y; by π^{xy} the derivative of $\pi(x, y; \lambda)$ with respect to x and to y; and treat similar expressions accordingly.

We then define the following class $\ddot{\Pi}^{1,1}(\lambda^*)$ of bidimensional poverty indices:

$$\ddot{\Pi}^{1,1}(\lambda^*) = \left\{ P(\lambda) \left| \begin{array}{l} \Lambda(\lambda) \subset \Lambda(\lambda^*) \\ \pi(x, y; \lambda) = 0, \text{ whenever } \lambda(x, y) = 0 \\ \pi^x \leq 0 \text{ and } \pi^y \leq 0 \; \forall x, y \\ \pi^{xy} \geq 0, \; \forall x, y. \end{array} \right. \right\} \tag{8}$$

The first line on the right of (8) defines the largest poverty set to which the poor must belong: the poverty set covered by the $P(\lambda)$ indices should lie within the maximal set $\Lambda(\lambda^*)$. The second line assumes that the poverty indices are continuous along the poverty frontier. This excludes the multidimensional poverty headcount, which is discontinuous at the poverty frontier. It implies among other things that small measurement errors should not have un-commensurate impacts on measured poverty. The third line of assumptions says that indices that are members of $\ddot{\Pi}^{1,1}$ are weakly decreasing in x and in y. This implies that an increase in either x or y cannot be bad for poverty reduction, Given the

[11]This differentiability assumption is made for expositional simplicity. It could be relaxed.
[12]The derivatives include the implicit effects of x and y on $\lambda(x, y)$.

interpretation of these variables, this would seem a natural assumption. For the indices to be non-degenerate, we must also have that $\pi^x < 0$, $\pi^y < 0$ and $\pi^{xy} > 0$ over some ranges of x and y. Note also that the inequalities in (8) are weak, which is different from the strong inequalities that are often found in the literature. This is consistent, however, with the way in which we will proceed to test dominance – we will test for strict ordering of the dominance surfaces, instead of the weak orderings often tested in the empirical literature.

The last line assumes that the marginal poverty benefit of an increase in either x or y decreases with the value of the other variable. Apart from our exclusion of the head-count, which we address below, this is the only assumption in (8) that implies debatable restrictions on the class $\ddot{\Pi}^{1,1}(\lambda^*)$, so we discuss it in some detail. Atkinson and Bourguignon (1982) and Bourguignon and Chakravarty (2002) refer to it as a property of non-decreasing poverty under a "correlation-increasing switch". To see what this implies, consider Figure 2. A first individual, initially located at a with an associated low x, sees his value of y also brought low by a movement from a to b. A second individual, who has more of x than the first individual, moves from c to d with a corresponding increase in his value of y by the same amount that the first individual sees his y fall. This is a "correlation-increasing switch". The marginal distributions of both x and y are unaffected by this switch. The joint distribution, $F(x, y)$, is however, increased by it. The correlation of deprivation and the incidence of multiple deprivation is higher after than before the switch.[13]

The $\pi^{xy} \geq 0$ assumption may also be understood as a "substitutability" assumption. The more someone has of x, the less is overall poverty deemed to be reduced if his value of y is increased. This assumption would seem to be ethically justified in many cases[14]. Governments, for instance, are often urged to care first for those who suffer from multiple deprivation, even though it may sometimes be more costly in budget resources to reach those individuals. An improvement in access to health services for those who are poorer in total expenditures would also seem socially more desirable than for those who are relatively better off in total expenditures.

But one can also think of other cases in which a strong complementarity in the produc-

[13]For further discussion of this concept, with examples, see Tsui (2002).

[14]See for instance Tsui (2002) for an advocate of this route.

tion of two dimensions of well-being might force us to reverse our assumption about pi^{xy} For instance, for a poverty analysis in the dimensions of education and nutritional status of children, there are production complementarities because better-nourished children learn better. If this complementarity is strong enough, it may overcome the usual ethical judgement that favors the multiply-deprived, so that overall poverty would decline by more if we were to transfer education from the poorly nourished to the better nourished, despite the fact that it increases the correlation of the two measures of wellbeing. Similarly, one might argue that human capital should be granted to those with a higher survival probability (because these assets would vanish following their death). Increasing the correlation of deprivations, and increasing the incidence of multiple deprivation, would then be good for poverty reduction. Bourguignon and Chakravarty (2002) derive dominance criteria for this second possibility. We do not pursue this avenue here. Our main reason is that this second approach would appear to limit drastically the scope for "poverty-frontier robust" orderings, in particular if robustness is sought over union/intersection/intermediate poverty-frontier definitions. [15] ***

Denote by $\Delta F = F_A - F_B$ the difference between a function F for A and for B. The class of indices defined in (8) then gives rise to the following $\ddot{\Pi}^{1,1}$ bi-dimensional dominance condition:

Theorem 1 **($\ddot{\Pi}^{1,1}$ poverty dominance)**

$$\Delta P(\lambda) > 0, \ \forall P(\lambda) \in \ddot{\Pi}^{1,1}(\lambda^*), \tag{9}$$

$$\textit{iff}\ \Delta P^{0,0}(x, y) > 0, \ \forall (x, y) \in \Lambda(\lambda^*). \tag{10}$$

Proof: See the appendix.

Condition (10) requires that the *bi-dimensional dominance surface* be higher for A than for B for all intersection poverty frontiers which lie in $\Lambda(\lambda^*)$. (The notation in (10) uses (x, y) and not (z_x, z_y) to stress that the comparison of the surfaces is to be done over an ***

[15]To see this, note that tests for classes of indices that obey the complementarity condition use dominance surfaces that are effectively defined by union poverty indices (see *e.g.* Bourguignon and Atkinson, 1982, and Bourguignon and Chakravarty, 2003). These tests thus include individuals lying beyond the intersection and intermediate poverty sets defined by the dominance surfaces (6) that are central to our proofs. The corresponding poverty tests with $\pi^{xy} \leq 0$ also suppose that there can be no welfare interaction between the two indicators of well-being whenever one's combination of indicators lies outside the inner intersection area (thus supposing that $\pi^{xy} = 0$ whenever $x > z_x$ or $y > z_y$ in Figure 1).

entire area, not only at some arbitrary combination of poverty lines.) If condition (10) is met, Theorem 1 says that poverty will be unambiguously higher in A than in B for all of the poverty indices that are members of the $\ddot{\Pi}^{1,1}(\lambda^*)$ class of poverty measures defined in Equation 8. Note that this allows for a wide area of poverty frontiers – all those such that $\Lambda(\lambda) \subset \Lambda(\lambda^*)$. This is an interesting and surprising feature of Theorem 1: even though the theorem is applicable to intersection, union, and intermediate poverty measures, one has only to check the intersection-like dominance surface in equation (6) over an appropriately defined domain to test for poverty dominance. Because of its generality, this result allows for powerful orderings of multidimensional poverty across A and B.

To see more clearly what Theorem 1 implies, return to Figure 1 and consider first an intersection poverty frontier for $\Lambda(\lambda^*)$. In this case, the relevant domain for the test would be a rectangle such as the one defined by the axes and the upper bounds (z_x, z_y). Thus, to establish a robust poverty comparison on this domain, we must check that A's dominance surface is above B's at every point over this rectangle. Note, however, that once this is established, we are assured of a robust poverty ordering not only at that precise intersection poverty frontier defined by (z_x, z_y), but also for any other poverty frontier which "fits" into this rectangle. These alternative poverty frontiers would include *all of the intermediate frontiers* (of the type of $\lambda_3(x, y) = 0$ in Figure 1) that could fit in the rectangle defined by $(0, 0)$ and (z_x, z_y). This is despite the fact that the bi-dimensional dominance surfaces are themselves only intersection poverty indices.

For a union poverty frontier λ^*, the test domain would be an L-shaped region defined in Figure 1 by $\Lambda(\lambda_2)$. Again, condition (10) requires that the dominance surface be higher for A than for B over all points within that region. If that is established, we are assured of a robust poverty ordering for all other union, intersection, or intermediate poverty frontiers and poverty sets which lie within that testing area. The extension to more general outer frontiers such as $\lambda_3(x, y) = 0$ in Figure 1 follows naturally.

2.3 Higher order dominance tests

For higher-order dominance, we may increase the order in one dimension or in both simultaneously. Either approach adds further assumptions on the effects of changes in either x or y on aggregate poverty, and thus limits the applicable class of poverty measures. These further assumptions are analogous to those found in the unidimensional dominance literature, and impose that indices react increasingly favorably to increases in living standards at the bottom of the distribution of well-being. The assumptions further require that the reactions of the indices to changes in one indicator be greater the lower the level of the other indicator of well-being.

To illustrate this, assume in addition to the above conditions for $\ddot{\Pi}^{1,1}$ that the first-order derivative $\pi^x(x, y; \lambda)$ is continuous whenever $\lambda(x, y) = 0$. Further suppose that equalizing transfers in x at a given value of y should, if enactable, weakly reduce aggregate poverty, and that this effect is decreasing in the value of y.[16] We then obtain the following class of bi-dimensional poverty indices:

***

$$\ddot{\Pi}^{2,1}(\lambda^*) = \left\{ P(\lambda) \left| \begin{array}{l} P(\lambda) \in \ddot{\Pi}^{1,1}(\lambda^*) \\ \pi^x(x, y; \lambda) = 0 \text{ whenever } \lambda(x, y) = 0 \\ \pi^{xx}(x, y; \lambda) \geq 0 \; \forall x, \\ \text{and } \pi^{xxy}(x, y; \lambda) \leq 0, \; \forall x, y. \end{array} \right. \right\} \tag{11}$$

This leads to the following dominance condition:

Theorem 2 **($\ddot{\Pi}^{2,1}$ poverty dominance)**

$$\begin{aligned} &\Delta P(\lambda) > 0, \; \forall P(\lambda) \in \ddot{\Pi}^{2,1}(\lambda^*) \\ &\textit{iff } \Delta P^{1,0}(x, y) > 0, \; \forall (x, y) \in \Lambda(\lambda^*). \end{aligned} \tag{12}$$

Proof: See appendix.

This tests simultaneous dominance of order 2 for x and of order 1 for y by checking whether the average poverty gap in x (given by $P^{1,0}(x, y)$), progressively cumulated in the dimension of y, is greater in A than in B, regardless of which intersection poverty frontier (x, y) is chosen within $\Lambda(\lambda^*)$. The ordering properties are analogous to those of Theorem 1.

Although it may not prove necessary, we can move to higher orders of dominance in the x dimension. The classes of poverty indices belonging to $\ddot{\Pi}^{s_x,1}(\lambda^*)$ become increasingly restricted as s_x increases. For $\ddot{\Pi}^{3,1}(\lambda^*)$ for instance, poverty indices must obey the principle of transfer sensitivity[17] in x, and react more to a favorable composite transfer the

[16]We say "if enactable" because, as a referee has pointed out, some dimensions of well-being such as children's heights are not literally transferable from one person to another.

[17]For a definition, see for instance Kakwani (1980) and Shorrocks and Foster (1987).

lower the value of y. Higher values of s_x imply compliance with higher-order principles of transfers [18].

In addition, we can simultaneously increase both s_x and s_y. The procedures, classes of poverty indices, and dominance relationships are analogous to those described above. For instance, the conditions for membership in $\ddot{\Pi}^{2,2}(\lambda)$ require that the poverty indices be convex in both x and y, and that they therefore obey the principle of transfers in both of these dimensions. They also require that this principle be stronger in one dimension of well-being the lower the level of the other dimension of well-being. Finally, they also impose that the second-order derivative in one dimension of well-being be convex in the level of the other indicator of well-being. The dominance condition then checks whether $P^{1,1}(x,y)$ is greater in A than in B for all $(x,y) \in \Lambda(\lambda^*)$.

2.4 Relevance of the methods

The methods that we propose above are more general than two other methods that researchers might use to consider poverty in multiple dimensions. One approach is to aggregate multiple dimensions of poverty into one univariate index, using arbitrary weights on each individual welfare measure. The best-known example is the Human Development Index (UNDP, 1990), which first constructs a scalar summary measure of each welfare variable, and then aggregates those measures into the HDI using arbitrary weights. An alternative approach in the same vein would be to sum each person's welfare variables, and then create a summary measure of that weighted sum for the population. In either case, the method arbitrarily reduces the problem of poverty comparison from many dimensions to one. Our proposed tests clearly generalize this rather restrictive approach.[19] ***

The second approach is to compare many indicators of well-being independently. Such univariate comparisons are also a special case of our approach. To see this, recall that the dominance surface $P^{s_x-1,s_y-1}(x,y)$ is cumulative in both dimensions. Hence, integrating out one dimension leaves the univariate dominance curve for the other dimension of well-being. In terms of Figure 1, the domain of separate univariate tests would be (for the y variable) a vertical line up to z_y at the maximum possible value of x, and (for

[18]For an illustration of this in the unidimensional setting, see Fishburn and Willig (1984).

[19]Dutta, Pattanaik, and Xu (2003) show that indices like the Human Development Index are very special and limited cases of more general multivariate welfare comparisons.

the x variable) a horizontal line up to z_x at the maximum possible value of y.

There are two ways in which our test could differ from this ”one-at-a-time” approach. First, it is possible that the univariate dominance curve for A lies above that for B at both $x = \infty$ and $y = \infty$ for the relevant range of poverty lines, but that A is not above B at one or more interior points in the test domain shown in Figure 1. In this case, the one-at-a-time approach would conclude that poverty is higher in A than B, but our bivariate approach would not. Because the bivariate approach checks the *joint distribution* of all indicators of well-being, it is thus able to show the correlation across such indicators, which is of ethical importance since it helps capture the joint incidence of deprivation in multiple dimensions. One-at-a-time analysis fails to do this.[20]

Alternatively , it is possible for the dominance surfaces to cross in the y dimension at $x = \infty$ and/or in the x dimension at $y = \infty$, but for A's surface to be above B's for a large area of interior points in the test domain. In this case, the one-at-a-time approach would not be able to establish a ranking of poverty, but our test would do so for intersection definitions of poverty and some intermediate definitions. Allowing for union definitions, however, would require including the margins of the surface in its test domain and would therefore not lead to a robust ordering of poverty.[21] ***

To visualize these methodological differences, consider Figure 3, which graphs a typical dominance surface. A larger ”hump” in the middle of the surface corresponds to a larger positive correlation between the two well-being variables. Also, the univariate dominance curve for one dimension is found at the upper extreme of the dominance surface's other dimension. On Figure 3, the univariate curve for the log of household expenditures lies on the extreme right of the surface, while that for the height-for-age z-score (to which we return below) is behind it.

When we make dominance comparisons, we test for the difference between two surfaces each like the one shown in Figure 3. Figure 4 depicts such a difference for the case in which one surface has highly correlated welfare variables while the second does

[20]We are grateful to a referee for pointing out that the one-at-a-time approach would provide necessary and sufficient conditions for a (smaller) class of union poverty measures that would allow *no* interaction between the indicators of well-being, namely, those for which $\pi^{xy} = 0$. See also the end of the proof of Theorem 1 in the appendix.

[21] In this case, one could first make the simpler univariate comparisons in each dimension, proceeding to the more difficult multivariate comparisons only if univariate dominance is found in each dimension. Crawford (1999) develops this approach.

not (the difference in these surfaces "bulges" in the middle). Although differences in the univariate dominance curves in both dimensions clearly cross the origin (at the extreme left and right of the figure), there is a significant interior section where the first surface is entirely above the second. Hence, there are intersection and intermediate poverty frontiers for which the first distribution has more poverty than the second. Conversely, we could think of shifting Figure 4 down such that the univariate differences are all negative. The first distribution would then dominate the second in both dimensions individually, but there would still be a section in the middle where the first surface would lie above the second. Thus, there would be no bivariate poverty dominance due to the first distribution showing too much incidence of multiple deprivation. We will give further examples of this in section 4.

2.5 Bounds to multidimensional dominance

Implementing the approaches to multidimensional poverty dominance developed above requires implicitly specifying a maximum poverty set $\Lambda(\lambda^*)$. Although there may be some intuitive feel that a very large set $\Lambda(\lambda^*)$ is not sensible, there is rarely reliable empirical evidence about what its precise value should be. Specifying this value *a priori* is thus necessarily subject to some degree of arbitrariness.

An alternative approach that gets around such arbitrariness is to estimate directly from the samples the maximum $\Lambda(\lambda^+)$ for which multidimensional poverty dominance holds in the sample. This *critical set* is delimited by a *critical poverty frontier*, since this will delimit the area of poverty frontiers which may not be exceeded for a robust multidimensional ordering of poverty to be possible. The researcher can then judge whether these critical sets and frontiers are sufficiently wide to justify a conclusion of poverty dominance.

To develop this idea further, assume that a critical poverty set exists in the two populations of multidimensional well-being being compared. Assume therefore that B initially dominates A but that their dominance surfaces eventually cross and that the ranking of the dominance surfaces is thus eventually reversed. Hence, for a given value of y, let $\zeta_x^+(y)$ then be the first crossing point[22] of the surfaces in the x dimension, with

[22]Note that $\zeta_x^+(y)$ will depend on the orders of dominance (s_x, s_y) considered, and should formally be written as $\zeta_x^+(y; s_x, s_y))$. For expositional simplicity, however, we do not make this dependence explicit.

$P_A^{\alpha_x,\alpha_y}(\zeta_x^+(y),y) = P_B^{\alpha_x,\alpha_y}(\zeta_x^+(y),y)$. Carrying out this exercise for a range of y leads to the estimation of a critical poverty frontier $\lambda^+(\zeta_x^+(y),y) \equiv 0$. By the results derived above, this procedure will provide an estimate of the space $\Lambda(\lambda^+)$ in which we can locate all of the possible poverty frontiers (union, intersection, or intermediate) for which there is necessarily more poverty in A than in B for all poverty indices that are members of $\ddot{\Pi}^{\alpha_x+1,\alpha_y+1}(\lambda^+)$. This procedure can be applied for any desired orders of bi-dimensional dominance $s_x = \alpha_x + 1$ and $s_y = \alpha_y + 1$, and can be generalized to more dimensions. Note that this poverty frontier $\lambda^+(x,y) = 0$ also locates the (x,y) frontier for which the bi-dimensional intersection FGT poverty indices would be exactly the same in the two distributions.

3 Estimation and inference

We now consider the estimation of the tools derived above for multidimensional poverty analysis. In this, we generalize to more than one dimension some of the results of Davidson and Duclos (2000).

3.1 Dominance surfaces

Suppose first that we have a random sample of N independently and identically distributed observations drawn from the joint distribution of x and y. We can write these observations of x^L and y^L, drawn from a population L, as (x_i^L, y_i^L), $i = 1, \ldots, N$. A natural estimator of the bidimensional dominance surfaces $P^{\alpha_x,\alpha_y}(z_x,z_y)$ is then: ***

$$\begin{aligned}
&\hat{P}_L^{\alpha_x,\alpha_y}(z_x,z_y) \\
&= \int_0^{z_y}\int_0^{z_x}(z_y-y)^{\alpha_y}(z_x-x)^{\alpha_x}\,d\hat{F}_L(x,y) \\
&= \frac{1}{N}\sum_{i=1}^N (z_y-y_i^L)^{\alpha_y}\,(z_x-x_i^L)^{\alpha_x} I(y_i^L \le z_y) I(x_i^L \le z_x) \\
&= \frac{1}{N}\sum_{i=1}^N (z_y-y_i^L)_+^{\alpha_y}\,(z_x-x_i^L)_+^{\alpha_x}
\end{aligned} \tag{13}$$

where $\hat{F}$ denotes the empirical joint distribution function, $I(\cdot)$ is an indicator function equal to 1 when its argument is true and 0 otherwise, and $x_+ = \max(0,x)$. For arbitrary α_x and α_y, (13) has the convenient property of being a simple sum of IID variables, even allowing for the fact that x and y and will generally be correlated across observations.

The following theorem allows us to perform statistical inference in the case in which we have a sample from each of two populations, A and B, that may or may not have been drawn independently from each other.

Theorem 3 *Let the joint population moments of order 2 of* $(z_y - y^A)_+^{\alpha_y}(z_x - x^A)_+^{\alpha_x}$ *and* $(z_y - y^B)_+^{\alpha_y}(z_x - x^B)_+^{\alpha_x}$ *be finite. Then* $N^{1/2}\left(\hat{P}_A^{\alpha_x,\alpha_y}(z_x,z_y) - P_A^{\alpha_x,\alpha_y}(z_x,z_y)\right)$ *and* $N^{1/2}\left(\hat{P}_B^{\alpha_x,\alpha_y}(z_x,z_y) - P_B^{\alpha_x,\alpha_y}(z_x,z_y)\right)$ *are asymptotically normal with mean zero, with asymptotic covariance structure given by* ($L, M = A, B$)*:*

$$\begin{array}{l} \lim_{N\to\infty} Ncov\left(\hat{P}_L^{\alpha_x,\alpha_y}(z_x,z_y), \hat{P}_M^{\alpha_x,\alpha_y}(z_x,z_y)\right) \\ = E\left((z_y - y^L)_+^{\alpha_y}(z_x - x^L)_+^{\alpha_x}(z_y - y^M)_+^{\alpha_y}(z_x - x^M)_+^{\alpha_x - 1}\right) \\ \quad -P_L^{\alpha_x,\alpha_y}(z_x,z_y)\, P_M^{\alpha_x,\alpha_y}(z_x,z_y) \end{array} . \tag{14}$$

Proof: See the appendix.

When the samples from the populations A and B are independent, the variance of each of $\hat{P}_A^{\alpha_x,\alpha_y}(z_x,z_y)$ and $\hat{P}_B^{\alpha_x,\alpha_y}(z_x,z_y)$ can be found by using (14) and by replacing N by N_A and N_B respectively. The covariance between the two estimators is then zero. The elements of the asymptotic covariance matrix can be estimated consistently using their sample equivalents.

3.2 Critical frontiers

To establish the sampling distribution of estimators of the critical frontier $\zeta_x^+(y)$, assume that within some bottom area $x \in [0, c_x]$ and at a given value of y, the population dominance surface for A lies above that for B, but that these surfaces cross (exactly) in the population at some higher critical point $\zeta_x^+(y)$. For a fixed value of y, a natural estimator $\hat{\zeta}_x^+(y)$ of the location of that point can be defined by the first point above y at which the sample ordering of the dominance surface changes. If the sample dominance surface for A were to lie always above that for B above y, then we could set $\hat{\zeta}_x^+(y)$ to an arbitrarily large value (denote it by z_x^+). Formally, $\hat{\zeta}_x^+(y)$ is then defined as[23]:

$$\hat{\zeta}_x^+(y) = \sup\left\{x | \Delta\hat{P}^{\alpha_x,\alpha_y}(x, y;) \geq 0 \text{ and } x \leq z_x^+\right\} \tag{15}$$

[23]Recall that $\Delta P = P_A - P_B$.

Applying this estimator over a range of y leads to an estimator of the critical frontier $\lambda^+(x, y) = 0$. Under these conditions, the sampling distribution of $\hat{\zeta}_x^+(y)$ is given by the following Theorem 4. For this, however, it is expositionally convenient to define an FGT index with a *negative* α_x as:

$$P^{-1,\alpha_y}(z_x, z_y) = \int_0^{z_y} (z_y - y)^{\alpha_y} f(y|x = z_x) dy \, f_x(z_x) \quad (16)$$

$$= E\left[(z_y - y)_+^{\alpha_y} |x = z_x\right] f_x(z_x) \quad (17)$$

where $f_x(z_x)$ is the marginal density of x at z_x and $f(y|x)$ is the conditional density of y at x. Both are assumed to exist. This leads to the following theorem.

Theorem 4 *Let the joint population moments of order 2 of* $\left(x^A\right)^{(\alpha_x)}\left(y^A\right)^{(\alpha_y)}$ *and* $\left(x^B\right)^{(\alpha_x)}\left(y^B\right)^{(\alpha_y-1)}$ *exist. If the samples from A and B are independent, assume that the ratio* $r = N_A/N_B$ *of their respective sample size tends to a constant as* N_A *and* N_B *tend to infinity. Under the conditions mentioned above (in particular, that* $\zeta_x^+(y)$ *exists in the population),* $N^{1/2}\left(\hat{\zeta}_x^+(y) - \zeta_x^+(y)\right)$ *is then asymptotically normal with mean zero, and its asymptotic variance is given by*

$$\begin{aligned}
& \lim_{N\to\infty} \mathit{var}\left(N^{1/2}\left(\hat{\zeta}_x^+(y) - \zeta_x^+(y)\right)\right) = \\
& \left[\delta \cdot \left(P_A^{\alpha_x-1,\alpha_y}(\zeta_x^+(y), y) - P_B^{\alpha_x-1,\alpha_y}(\zeta_x^+(y), y)\right)\right]^{-2} \\
\times \; & \left[\mathrm{var}\left((y - y^A)_+^{\alpha_y}\left(\zeta_x^+(y) - x^A\right)_+^{\alpha_x}\right) + \mathrm{var}\left((y - y^B)_+^{\alpha_y}\left(\zeta_x^+(y) - x^B\right)_+^{\alpha_x}\right)\right. \\
& \left. -2\,\mathrm{cov}\left((y - y^A)_+^{\alpha_y}\left(\zeta_x^+(y) - x^A\right)_+^{\alpha_x}, (y - y^B)_+^{\alpha_y}\left(\zeta_x^+(y) - x^B\right)_+^{\alpha_x}\right)\right] \quad (18)
\end{aligned}$$

when the samples are dependent, and by

$$\begin{aligned}
& \lim_{N_A\to\infty} \mathit{var}\left(N_A^{1/2}\left(\hat{\zeta}_x^+(y) - \zeta_x^+(y)\right)\right) = \\
& \left[\delta \cdot \left(P_A^{\alpha_x-1,\alpha_y}(\zeta_x^+(y), y) - P_B^{\alpha_x-1,\alpha_y}(\zeta_x^+(y), y)\right)\right]^{-2} \\
\times \; & \left[\mathrm{var}\left((y - y^A)_+^{\alpha_y}\left(\zeta_x^+(y) - x^A\right)_+^{\alpha_x}\right) + r\,\mathrm{var}\left((y - y^B)_+^{\alpha_y}\left(\zeta_x^+(y) - x^B\right)_+^{\alpha_x}\right)\right] \quad (19)
\end{aligned}$$

when the samples are independent, and by setting $\delta = \alpha_x$ *when* $\alpha_x > 0$*, and* $\delta = 1$ *when* $\alpha_x = 0$.

Proof: See appendix.

4 Examples

As a first example, consider the question: are rural people poorer than the urban ones in Viet Nam? Many studies, of Viet Nam and elsewhere, find that people living in rural areas tend to be poorer when judged by expenditures or income alone. However, it is possible that people are better nourished in rural than urban areas, *ceteris paribus*, because they have tastes for foods that provide nutrients at a lower cost, or because unit prices of comparable food commodities are lower. In such cases, including an indicator of nutritional status may change the relative well-being of rural and urban residents. To test this, we measure welfare in two dimensions: *per capita* household expenditures and nutritional status, as measured by a child's gender- and age-standardized height, transformed into standard deviations from the median of a healthy population, known as z-scores. Stunted growth in children is widely used as an indicator of malnutrition and poor health. The sample comes from the Viet Nam Living Standards Measurement Survey carried out in 1993[24]. This is a nationally representative household survey that collected detailed expenditure and anthropometric data. The latter, however, are available only for children younger than 60 months, so our sample is actually these children only, rather than all of the members of the households interviewed.

The test described in equation (10) requires comparison of the two dominance surfaces of urban and rural children in Viet Nam: the difference between those two surfaces is shown in Figure 5 for $s_x = s_y = 1$. The y axis measures the height-for-age z-score (stunting); the x axis measures the *per capita* expenditures for the child's household; and the z axis measures the cumulative proportion of children that fall below the points defined in the (x, y) domain. The poorest children are in the front left-hand corner of the graph. If the rural dominance surface is above the urban surface over the relevant area of poverty frontiers (values of per capita expenditures and stunting), poverty is higher (more people are below the given well-being levels in each dimension) in rural areas. This conclusion is then robust to the choice of poverty indices in the class $\ddot{\Pi}^{1,1}(\lambda^*)$.

In theory, we should test over the entire area defined by $\Lambda(\lambda^*)$, but it is more practical

[24]Information on the LSMS surveys is available in Grosh and Glewwe (1998). Information about the specific survey that we use is available at the LSMS website: http://www.worldbank.org/html/prdph/lsms/country/vn93/vn93bid.pdf.

to choose a grid of points in the (x, y) domain over which to test[25] Here we use a grid that is 20x20, spread evenly over the *entire* domain of the log of per capita expenditures and the height-for-age z-score. Following Howes (1996), we test for a significant difference in the dominance surface at each point of the grid, and reject the null of non-dominance of A by B only if all of the test statistics (t-statistics) have the right sign and are significantly different from 0.

Figure 5 indicates clearly that, over almost the entire range of expenditures and stunting, rural children are poorer than urban. Table 1 shows whether these statements are statistically significantly at the 5% level. A negative sign indicates that the urban dominance surface is significantly below the rural one, a positive sign indicates the opposite, and a zero indicates that the difference is not statistically significant. The negative differences are statistically significant for any reasonable pair of poverty lines (except at the very bottom right of Table 1. Hence, by Theorem 1, the conclusion that rural children are poorer than urban ones is valid for almost any intersection, union or intermediate poverty frontier.

Our second example tests for first-order poverty dominance in three dimensions. We ask whether poverty declined in Ghana between 1993 and 1998, using data from the Demographic and Health Surveys. The three welfare variables that we consider are for children under five years old: their survival probability, their height-for-age z-score (stunting), and an index of their household's assets.[26] We compare dominance surfaces for these three measures in 1993 and 1998, the two years for which DHS data exist. While we cannot graph the resulting four-dimensional surface, Figure 6 summarizes the results of the statistical test. We use a 20x20x20 grid of test points, and each horizontal layer in Figure 6 is similar to Table 1 in the previous example.[27] A light gray plus "+" sign indicates that the 1998 surface is significantly above the 1993 surface; a circle indicates that the 1998 surface is significantly below the 1993 surface; and a solid black point indicates that they are statistically indistinguishable at the five-percent significance level. It is clear from the figure that there is no robust poverty dominance result. Over some of the domain, poverty

[25]The computation of these statistics, as well as of their sampling distribution, can be done using the software DAD (Duclos, Araar and Fortin (2002)), freely available at www.mimap.ecn.ulaval.ca.

[26]Information on the estimation of survival probabilities and the asset index may be found in Sahn, Stifel, and Younger (1999) and Sahn and Stifel (2000).

[27]We have excluded some of the horizontal layers to make the graph more legible.

does seem to have declined between 1993 and 1998. But in significant areas, particularly for low values in the asset dimension, the reverse is true.

In addition to showing that our tests are possible in more than two dimensions, this example shows the importance of checking for the robustness of poverty comparisons using tests such as those that we employ. For the intersection headcount, shown by points on the dominance surfaces, a judicious choice of the poverty lines could lead one to conclude that poverty worsened, improved, or did not change, depending on the specific choice. None of these results would be robust, but any might seem plausible if it appeared on its own.

The next two examples highlight the difference between using bivariate dominance tests *vs.* one-at-a-time univariate tests on the same variables. Table 2 gives the results for tests of the differences in the dominance surfaces for stunting and child survival probability in Cameroon and Madagascar. The data come from the 1997 Demographic Health Surveys (DHS) in those countries.[28] The "one-at-a-time" dominance curves are given in the last row of the table (for survival probability) and in the last column (for stunting). These univariate comparisons would conclude that poverty is worse in Madagascar than in Cameroon, whether measured by stunting or survival probability. Nevertheless, the bivariate comparison shows several internal points where the surfaces are not significantly different, including two where the point estimate of the difference is in fact positive. So our method would not come to the same conclusion, finding instead that there is no statistically significant, first-order poverty ordering of these two populations.

Table 3 shows the other possibility for different conclusions. These results are also for tests of the differences between first-order dominance surfaces for stunting and child survival probability, in Colombia and the Dominican Republic, and come from the DHS surveys for those countries, carried out in 1995 and 1996 respectively. In this case, there is dominance on one margin (for survival probability), but not the other (stunting), so the one-at-a-time approach would not find poverty to be necessarily lower in one population than the other. However, Colombia's dominance surface is significantly below the Dominican Republic's over a very large range of the interior points, suggesting that under an intersection definition of poverty, and several intermediate ones as well, poverty was

[28]Information on these surveys is available at http://www.measuredhs.com. The child survival probability is estimated as in Sahn, Stifel, and Younger (1999).

robustly lower in Colombia than in the Dominican Republic.

Turning our attention now to the analysis of critical poverty frontiers, Figure 7 shows two such frontiers, for the $\Pi^{1,1}$ and $\Pi^{2,2}$ classes of poverty measures, respectively, using children's height-for-age z-score and the log of their households' per capita expenditures as measures of well-being.[29] The data are from the 1999 Uganda National Household Survey, and the comparison is between rural residents in the Eastern region and urban residents in the Northern region. Up to these critical frontiers, poverty is lower in rural Eastern Uganda than it is in urban Northern Uganda for all poverty measures in the respective class. Note that the critical frontier is close to the origin for $\Pi^{1,1}$, so that relatively few poverty frontiers fit within the critical frontier.[30] For $\Pi^{2,2}$, however, the critical frontier extends much farther, so that a rather large set of intersection and intermediate poverty measures in this class conclude that poverty is lower in rural Eastern *vs.* urban Northern Uganda. Of course, the price of this result is to be valid only for a smaller class $\Pi^{2,2}$ of poverty measures, a class that requires *inter alia* poverty indices to fall following equalizing transfers in dimensions x or y – see Section 2.3. ***

5 Conclusion

In this paper, we have shown that it is possible to make very general poverty comparisons for multiple dimensions of well-being. These comparisons have several attractive features:

1. In the spirit of the stochastic dominance literature, they can be tested for robustness over broad classes of poverty aggregation procedures, both within a given dimension of well-being and across the different dimensions of interest.

2. A special consideration for the multivariate case is whether poverty is defined as the intersection or union of poverty in each dimension. The methods that we describe are valid for both, as well as for any choice of intermediate cases.

3. The poverty comparisons can be tested for robustness over a broad area of poverty frontiers. Alternatively, one can estimate a critical poverty frontier up to which multidimensional poverty dominance necessarily holds.

[29]The figure actually shows the critical frontier minus two standard deviations.
[30]As a point of reference, the median z-score is -1.22 and the median log of per capita expenditures is 8.52.

4. The comparisons involve statistical tests that make use of the sampling distribution of multidimensional dominance surfaces.

The importance of these methods rests on two considerations. The first is ethical: there is widespread acknowledgement that well-being and poverty are multidimensional in nature. We take this as given, based either on Sen's philosophical arguments that poverty should be viewed in terms of multidimensional capabilities and functionings, or on the realistic recognition that a person's well-being has dimensions that cannot be purchased and that transcend levels of income. The second consideration is practical: to what extent will multidimensional comparisons differ from unidimensional ones? Given the relatively weak correlations that are often observed between income and other welfare variables, it should not be surprising to find cases where poverty comparisons in multiple dimensions differ from comparisons in only one of those dimensions, something that we found repeatedly in preparing the examples for this paper.

More generally, we have shown that our multidimensional comparisons can also differ from univariate comparisons in each individual dimension in two ways. One could find dominance in each dimension separately if the margins of the dominance surfaces differ in each dimension, but may not find multidimensional dominance if the surfaces cross in the surfaces' interiors. This draws attention to the importance of capturing the incidence of multiple deprivation. Alternatively, the one-dimensional dominance curves may cross, ruling out univariate dominance, but the interiors of the multidimensional surfaces may be uniformly different, allowing multivariate dominance for intersection or intermediate poverty definitions. The empirical importance of these two possibilities remains to be firmly established. Our admittedly limited experience based on comparisons of the DHS surveys is that the first is rare, while the second is fairly common. But a firm sense of the importance of our more general methods must await further practice with other samples and other variables.

6 Appendix

Proof of Theorem 1. ∎

Denote the points on the outer poverty frontier $\lambda^*(x, y) = 0$ as $z_x(y)$ for a point above y and $z_y(x)$ for a point above x. The derivative conditions in (1) imply that $z_x^{(1)}(y) \leq$

0 and $z_y^{(1)}(x) \leq 0$, where the superscript (1) indicates the first-order derivative of the function with respect to its argument. Note that the inverse of $z_x(y)$ is simply $z_y(x)$: $x \equiv z_x(z_y(x))$.

We then proceed by first integrating equation (3) by parts with respect to x, over an interval of y ranging from 0 to z_y. z_y can extend to infinity if needed, *e.g.,* for a union poverty frontier. 0 stands for a lower bound that could also be set arbitrarily low, thus allowing for negative values of well-being indicators. This gives:

$$
\begin{aligned}
P(z_x(y), z_y) &= \int_0^{z_y} \left[\pi(x, y; \lambda^*)\, F(x|y)\right]\Big|_0^{z_x(y)} f(y)\, dy \\
&- \int_0^{z_y} \int_0^{z_x(y)} \pi^x(x, y; \lambda^*)\, F(x|y) f(y)\, dx\, dy. \qquad (20)
\end{aligned}
$$

The first term on the right-hand-side of (20) is zero since $F(x = 0|y) = 0$ and since we assumed that $\pi(z_x(y), y; \lambda^*) = 0$. Hence, it is here that the continuity assumption at the poverty frontiers is technically important. To integrate by parts with respect to y the second term, define a general function $K(y) = \int_0^{g(y)} h(x, y)\, l(x, y) dx$ and note that:

$$
\begin{aligned}
\frac{dK(y)}{dy} &= g^{(1)}(y) h(g(y), y)\, l(g(y), y) \\
&+ \int_0^{g(y)} \frac{\partial h(x, y)}{\partial y}\, l(x, y) dx \\
&+ \int_0^{g(y)} h(x, y)\, \frac{\partial l(x, y)}{\partial y} dx. \qquad (21)
\end{aligned}
$$

Reordering (21) and integrating it from 0 to c, we find:

$$
\begin{aligned}
&- \int_0^c \int_0^{g(y)} h(x, y)\, \frac{\partial l(x, y)}{\partial y} dx dy \\
&= -K(c) + K(0) + \int_0^c g^{(1)}(y) h(g(y), y)\, l(g(y), y) dy \\
&+ \int_0^c \int_0^{g(y)} \frac{\partial h(x, y)}{\partial y}\, l(x, y) dx dy. \qquad (22)
\end{aligned}
$$

Now replace in (22) c by z_y, $g(y)$ by $z_x(y)$, $h(x, y)$ by $\pi^x(x, y; \lambda^*)$, $l(x, y)$ by $F(x, y)$ and $K(y)$ by its definition $K(y) = \int_0^{g(y)} h(x, y)\, l(x, y) dx$. This gives:

$$
\begin{aligned}
P(z_x(y), z_y) &= -\int_0^{z_x(z_y)} \pi^x(x, z_y; \lambda^*)\, P^{0,0}(x, z_y)\, dx & (23)\\
&+ \int_0^{z_y} z_x^{(1)}(y)\, \pi^x(z_x(y), y; \lambda^*)\, P^{0,0}(z_x(y), y)\, dy & (24)\\
&+ \int_0^{z_y} \int_0^{z_x(y)} \pi^{xy}(x, y; \lambda^*)\, P^{0,0}(x, y)\, dx\, dy. & (25)
\end{aligned}
$$

For the sufficiency of condition (10), recall that $z_x^{(1)}(y) \le 0$, $\pi^x \le 0$, and $\pi^{xy} \ge 0$, with strict inequalities for either of these inequalities over at least some inner ranges of x and y. Hence, if $\Delta P^{0,0}(x, y) > 0$, for all $y \in [0, z_y]$ and for all $x \in [0, z_x(y)]$ (that is, for all $(x, y) \in \Lambda(\lambda^*)$), then it must be that $\Delta P(\lambda^*) > 0$ for all of the indices that use the poverty set $\Lambda(\lambda^*)$ and that obey the first two lines of conditions in (8). But note that for other poverty sets $\Lambda(\lambda) \subset \Lambda(\lambda^*)$, the relevant sufficient conditions are only a subset of those for $\Lambda(\lambda^*)$. The sufficiency part of Theorem 1 thus follows.

For the necessity part, assume that $\Delta P^{0,0}(x, y) \le 0$ for an area defined over $x \in [c_x^-, c_x^+]$ and $y \in [c_y^-, c_y^+]$, with $c_x^+ \le z_y$ and $c_x^+ \le z_x(y)$. Then any of the poverty indices in $\ddot{\Pi}^{1,1}(\lambda^*)$ for which $\pi^{xy} < 0$ over that area, $\pi^{xy} = 0$ outside that area, and for which $\pi^x(x, z_y; \lambda^*) = \pi^x(z_x(y), y; \lambda^*) = 0$, will indicate that $\Delta P < 0$. Condition (10) is thus also a necessary condition for the ordering specified in Theorem 1. (Note, however, that necessary and sufficient conditions for a subclass of union poverty indices with $\pi^{xy} = 0$ would only involve the marginal or univariate distributions – this can be seen by inspection of (25).). ***

∎

Proof of Theorem 2. ∎

Integrating (25) once more by parts with respect to x, and imposing the continuity conditions characterizing the indices $\ddot{\Pi}^{2,1}(\lambda^*)$ in (11), we find:

$$
\begin{aligned}
P(\lambda^*) &= \int_0^{z_x(z_y)} \pi^{xx}(x, z_y; \lambda^*) P^{1,0}(x, z_y) dx\\
&+ \int_0^{z_y} \pi^{xy}(z_x(y), y; \lambda^*) P^{1,0}(z_x(y), y) dy\\
&- \int_0^{z_y} \int_0^{z_x(y)} \pi^{xxy}(x, y; \lambda^*) P^{1,0}(x, y) dx\, dy. & (26)
\end{aligned}
$$

The rest of the proof is as for Theorem 1.

∎

Proof of Theorem 3.

For each distribution, the existence of the appropriate population moments of order 1 lets us apply the law of large numbers to (13), thus showing that $\hat{P}_K^{\alpha_x,\alpha_y}(z_x, z_y)$ is a consistent estimator of $P_K^{\alpha_x,\alpha_y}(z_x, z_y)$. Given also the existence of the population moments of order 2, the central limit theorem shows that the estimator in (13) is root-N consistent and asymptotically normal with asymptotic covariance matrix given by (14). When the samples are dependent, the covariance between the estimator for A and for B is also provided by (14).

Theorem 3 thus provides the formula needed to estimate the sampling variability of any point on the dominance surfaces and for any choice of multidimensional poverty lines in the multidimensional FGT poverty indices. Extension of the result of Theorem 3 to any additive multidimensional poverty indices is straightforward, and simply requires substituting in (14) the relevant functions $\pi(x, y; \lambda)$ for $((z_y - y)_+^{\alpha_y} (z_x - x)_+^{\alpha_x}$.

∎

Proof of Theorem 4.

The proof can be established along the lines of the proof of Theorem 3 in Davidson and Duclos (2000). To see this, note that the conditions of Theorem 4 assume that the appropriate joint population moments exist, and that the critical frontier $\zeta_x^+(y)$ also exists in the population. Furthermore, since this frontier is assumed to be where the dominance surfaces *exactly* cross, we have that $\Delta \left(\partial P^{s_x-1,s_y-1}(x,y)/\partial x\right)\big|_{x=\zeta_x^+(y)} < 0$. Note that this derivative is given by $\delta \cdot \left(P_A^{s_x-2,s_y-1}(\zeta_x^+(y), y) - P_B^{s_x-2,s_y-1}(\zeta_x^+(y), y)\right)$, with $\delta = s_x - 1$ when $s_x > 1$, and $\delta = 1$ when $s_x = 1$. When $s_x = 1$, we also have $P_L^{-1,s_y-1}(\zeta_x^+(y), y) = E\left[(y - y^L)_+^{(s_y-1)} | x = \zeta_x^+(y)\right] f_x(\zeta_x^+(y)), L = A, B$.

Again, the elements of the asymptotic covariance matrix can be estimated consistently by simply using their sample estimates. Estimating $P^{s_x-2,s_y-1}(\zeta_x^+(y), y)$ is also easily done when $s_x > 1$. Estimating $E\left[(y - y^L)_+^{(s_y-1)} | x = \zeta_x^+(y)\right] f_x(\zeta_x^+(y))$ is slightly more complicated, but can be done consistently using non-parametric regression procedures. In particular, we use in the illustration a Gaussian kernel, $K(u) = (2\pi)^{-0.5} \exp^{-0.5u^2}$, and estimate $E\left[(y - y^L)_+^{(s_y-1)} \middle| x = \zeta_x^+(y)\right] f_x(\zeta_x^+(y))$ as:

$$(nh)^{-1} \sum_{i=1}^{n} K\left(\frac{\zeta_x^+(y) - x_i^L}{h}\right) \left(y - y_i^L\right)_+^{(s_y-1)}. \tag{27}$$

References

[1] Atkinson, A.B. (1987). "On the Measurement of Poverty", *Econometrica*, **55**, 749–764.

[2] Atkinson, A.B. (1991). "Measuring Poverty and Differences in Family Composition", *Economica*, **59**, 1–16.

[3] Atkinson, A.B. (2002). "Multidimensional Deprivation: Constrasting Social Welfare and Counting Approaches", mimeo, Nuffield College, Oxford University.

[4] Atkinson, A.B. and F. Bourguignon (1982). "The Comparison of Multi-Dimensional Distributions of Economic Status", chapter 2 in *Social Justice and Public Policy*, Harvester Wheatsheaf, London.

[5] Atkinson, A.B. and F. Bourguignon (1987). "Income Distribution and Differences in Needs", in G.R. Feiwel, ed., *Arrow and the foundations of the theory of economic policy*, New York Press, New York, 350–70.

[6] Bellman, R. (1961). *Adaptive Control Processes: A Guided Tour*, Princeton University Press.

[7] Bourguignon, F. (1989). "Family Size and Social Utility: Income Distribution Dominance Criteria", *Journal of Econometrics*, **42**, 67–80.

[8] Bourguignon, F., and S. R. Chakravarty (2003). "The measurement of multidimensional poverty", *Journal of Economic Inequality*), **1**, 25–49.

[9] Bourguignon, F. and S.R. Chakravarty, (2002). "Multi-dimensional poverty orderings", DELTA, Paris.

[10] Bourguignon, F. and G. Fields (1997). "Discontinuous Losses from Poverty, Generalized P Measures, and Optimal Transfers to the Poor", *Journal of Public Economics*, **63**,155–175.

[11] Chakravarty, S.R. (1983). "A New Index of Poverty", *Mathematical Social Sciences*, **6**, 307–313.

[12] Chambaz, C. and E. Maurin (1998). ”Atkinson and Bourguignon Dominance Criteria: Extended and Applied to the Measurement of Poverty in France”, *Review of Income and Wealth*, **44**, 77–124.

[13] Crawford, Ian A. (1999). ”Nonparametric Tests of Stochastic Dominance in Bivariate Distributions, with an Application to UK”, University College London Discussion Papers in Economics 99/07.

[14] Davidson, R. and J.-Y. Duclos (2000). “Statistical Inference for Stochastic Dominance and the for the Measurement of Poverty and Inequality”, *Econometrica*, **68**, 1435–1465.

[15] Duclos, J.-Y., Araar, A. and C. Fortin (2002), ”DAD: A Software for Distributive Analysis / Analyse distributive”, MIMAP, International Research Centre, Government of Canada and CIRPÉE, Universit Laval (www.mimap.ecn.ulaval.ca).

[16] Duclos, J.-Y. and P. Makdissi (2000). “Sequential Stochastic Dominance and the Robustness of Poverty Orderings”, Cahier de recherche, Département d’économique, Université Laval.

[17] Dutta, I., P.K. Pattanaik, and Y. Xu (2003). “On Measuring Deprivation and the Standard of Living in a Multidimensional Framework on the Basis of Aggregate Data,” *Economica*, **70**, 197–221.

[18] Fishburn, P.C. and R.D. Willig (1984). ”Transfer Principles in Income Redistribution”, *Journal of Public Economics*, **25**, 323–328.

[19] Foster, J.E., (1984). ”On Economic Poverty: A Survey of Aggregate Measures”, in R.L. Basmann and G.F. Rhodes, eds., *Advances in Econometrics*,**3**, Connecticut: JAI Press, p. 215-251.

[20] Foster, J.E., J. Greer and E. Thorbecke (1984). “A Class of Decomposable Poverty Measures”, *Econometrica*, **52 (3)**, 761–776.

[21] Foster, J.E. and A.F. Shorrocks (1988a). “Poverty Orderings”, *Econometrica*, **56**, 173–177.

[22] Foster, J.E. and A.F. Shorrocks (1988b). “Poverty Orderings and Welfare Dominance”, *Social Choice Welfare*, **5**, 179–198.

[23] Foster, J.E. and A.F. Shorrocks (1988c). “Inequality and Poverty Orderings”, *European Economic Review*, **32**, 654–662.

[24] Government of Romania (1994). Romania Integrated Household Survey.

[25] Grosh, B. and P. Glewwe (1998). “Data Watch: The World Bank’s Living Standards Measurement Study Household Surveys”, *Journal of Economic Perspectives*, **12**, 187– 96.

[26] Howes, S. (1996). “A new test for inferring dominance from sample data”, draft.

[27] Jenkins, S.P. and P.J. Lambert (1993). “Ranking Income Distributions when Needs Differ”, *Review of Income and Wealth*, **39**, 337–356.

[28] Kakwani, N.C. (1980). “On a Class of Poverty Measures” *Econometrica*, **48**, 437–446.

[29] Sahn, David E. and David C. Stifel (2000). ”Poverty Comparisons Over Time and Across Countries in Africa,” *World Development*, Vol. 28, Number 12.

[30] Sahn, David E., David Stifel, and Stephen Younger (1999). ”Inter-temporal Changes in Welfare: Preliminary Results from Nine African Countries,” CFNPP Working Paper #94.

[31] Sen, A.K., (1979). ”Personal utilities and public judgment: Or what’s wrong with welfare economics?” *Economic Journal* Vol. 89 Number 355, 537-558.

[32] Sen, A.K. (1985), *Commodities and Capabilities*, North-Holland, Amsterdam.

[33] Sen, A.K., (1987). ”The Standard of Living: Lecture II, Lives and capabilities,” in: Hawthorn, G. (ed.), *The Standard of Living*. Cambridge University Press, Cambridge, pp. 20-38.

[34] Shorrocks, A.F., and J. Foster (1987). ”Transfer Sensitive Inequality Measures”, *Review of Economic Studies*, **LIV**, 485–497.

[35] Tsui, K, (2002). ”Multidimensional poverty indices”, *Social Choice and Welfare*, **19**, 69–93.

[36] United Nations Development Program (1990). *Human Development Report*. New York: Oxford University Press.

[37] Watts, H.W. (1968), "An Economic Definition of Poverty", in D.P. Moynihan (ed.), *On Understanding Poverty*, New York: Basic Books.

Figure 1: Union and intersection poverty indices

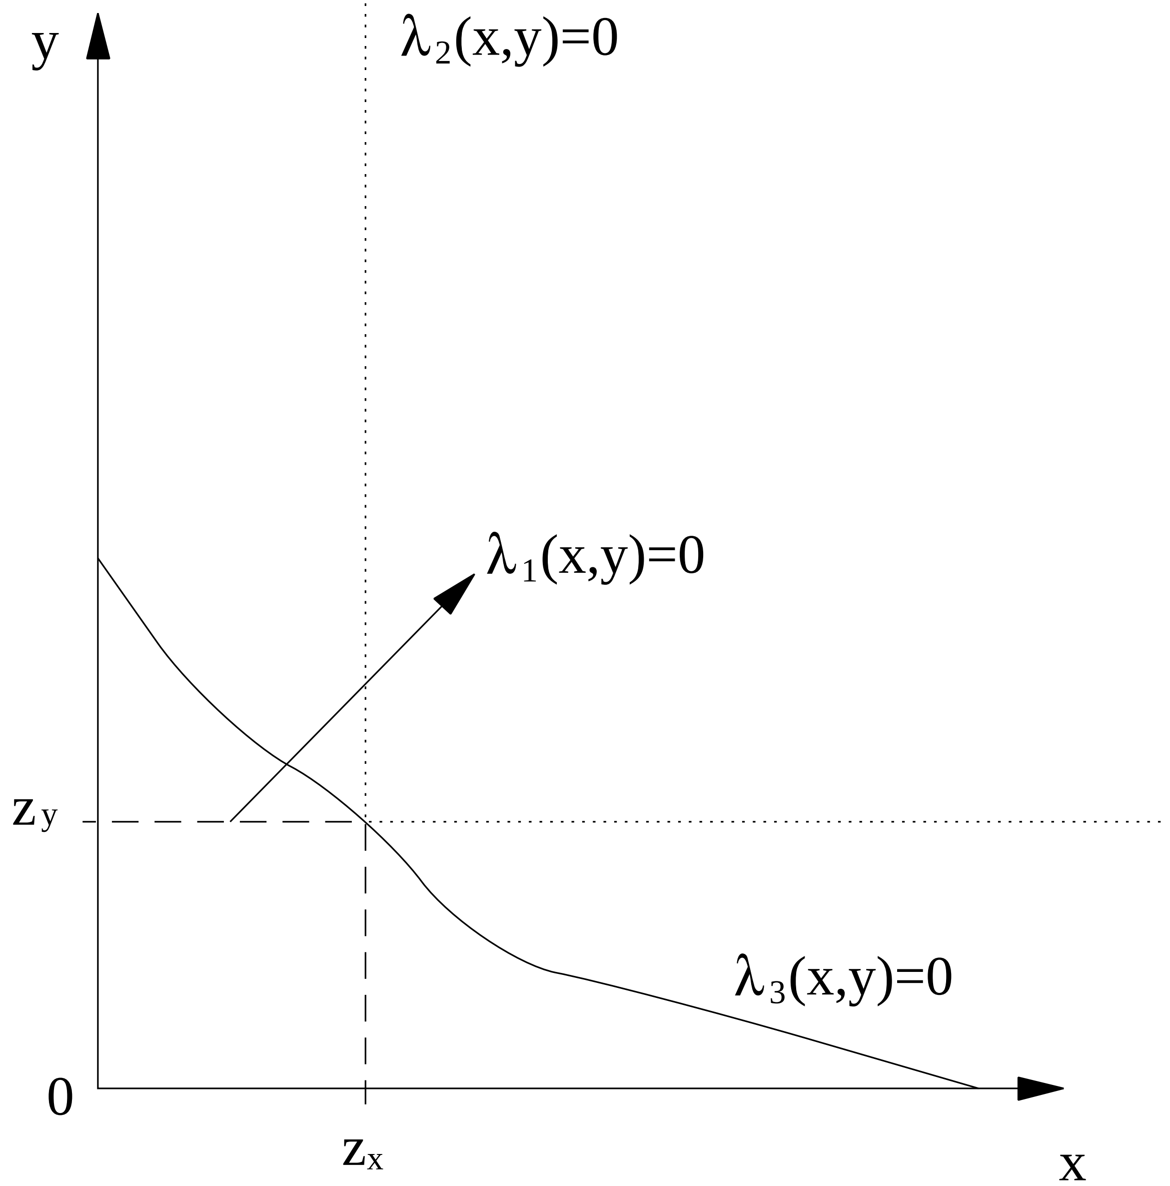

Figure 2: A correlation-increasing switch

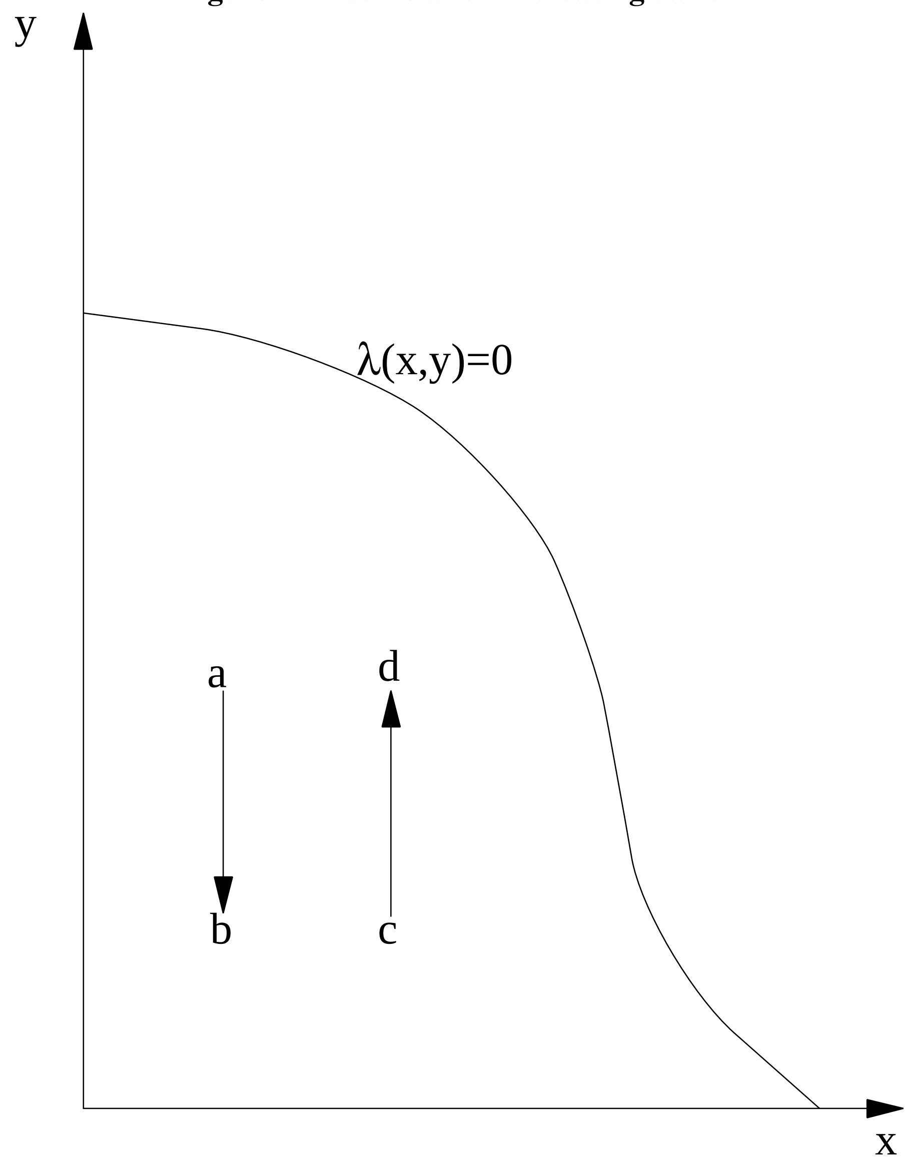

Figure 3: Dominance surface for Ghanaian children, 1989

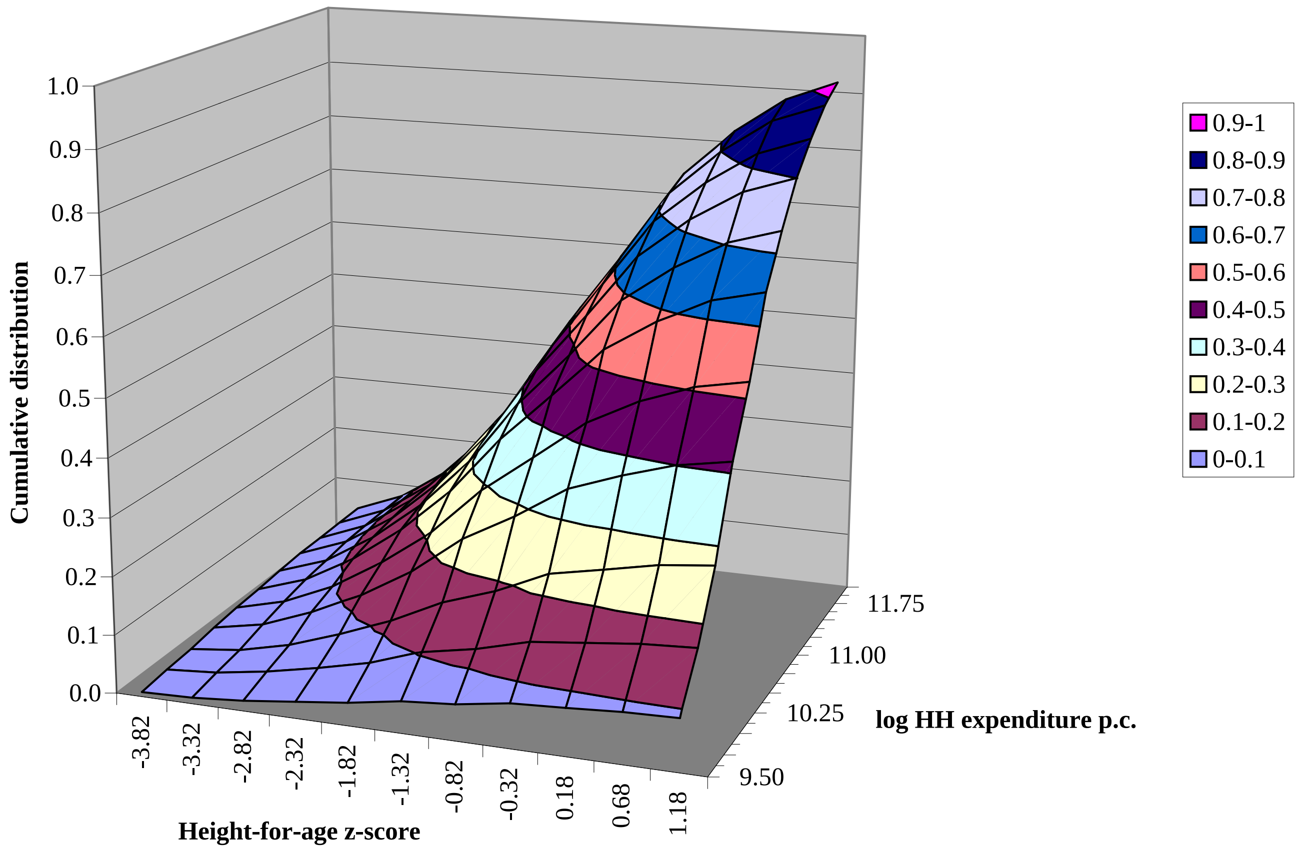

Figure 4: Example of difference in dominance surfaces, intersection dominance without marginal dominance

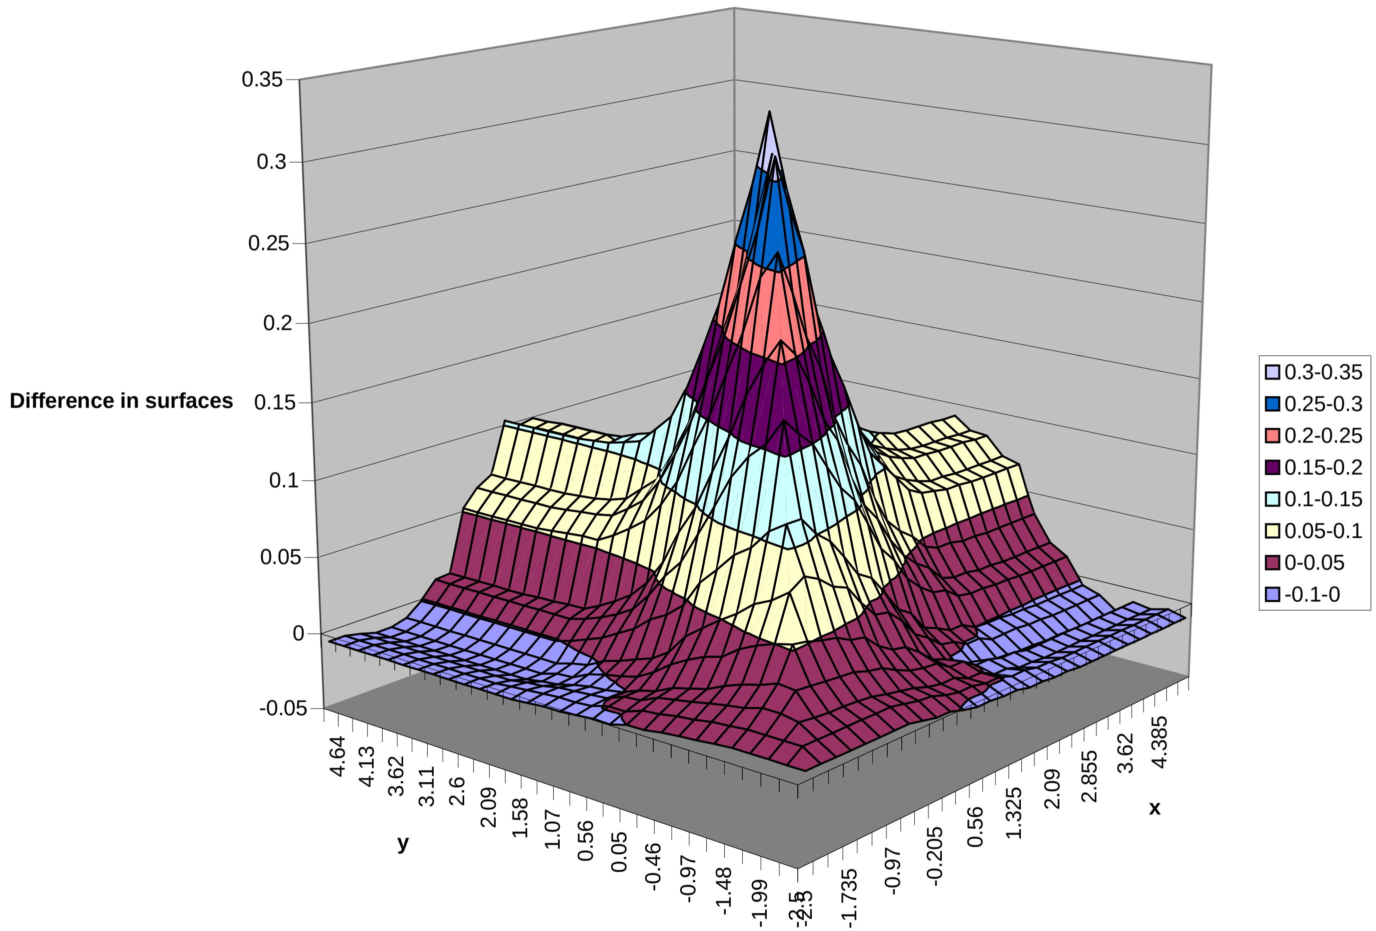

Figure 5: Urban minus Rural Dominance Surface for Viet Nam

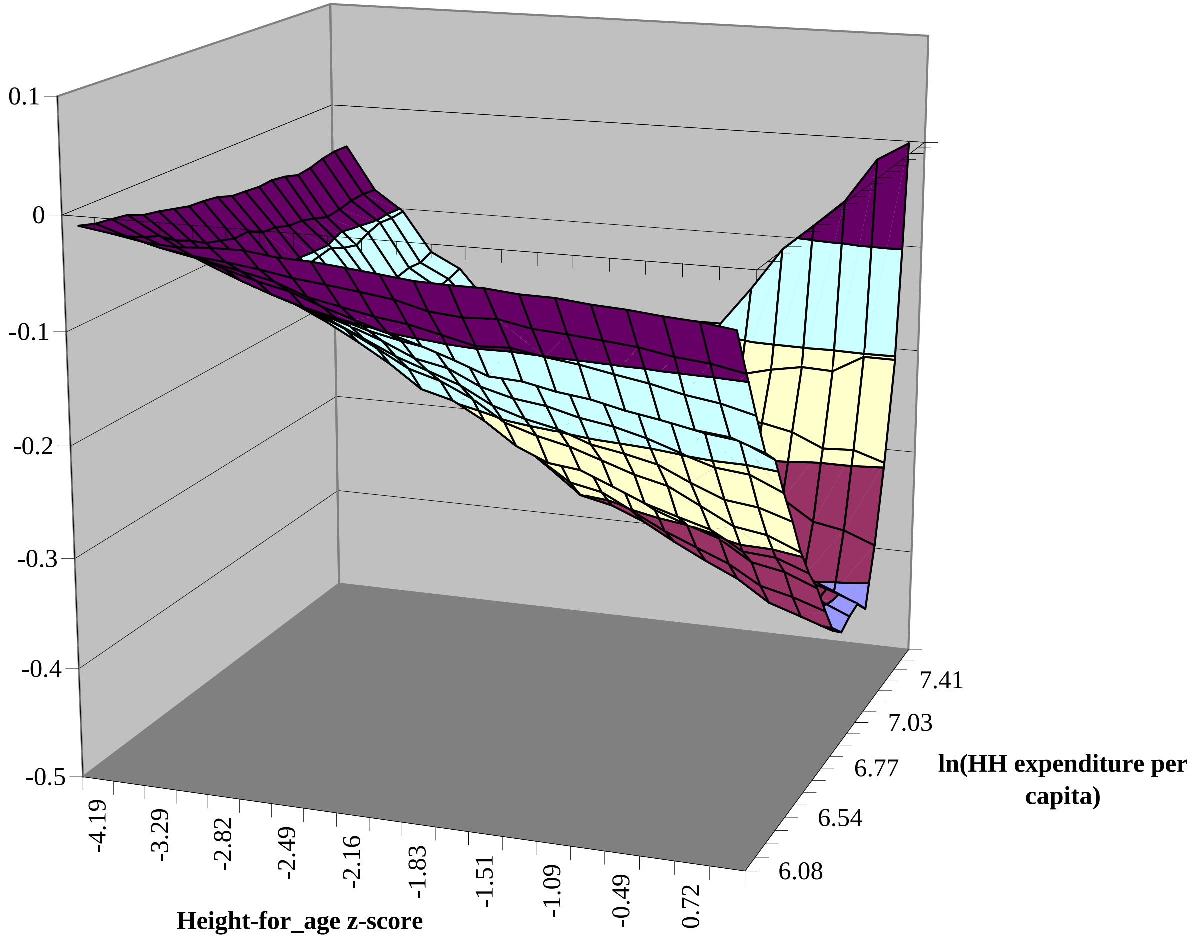

Figure 6: Test results for difference between 1993 and 1998 first-order dominance surfaces for Ghanaian children

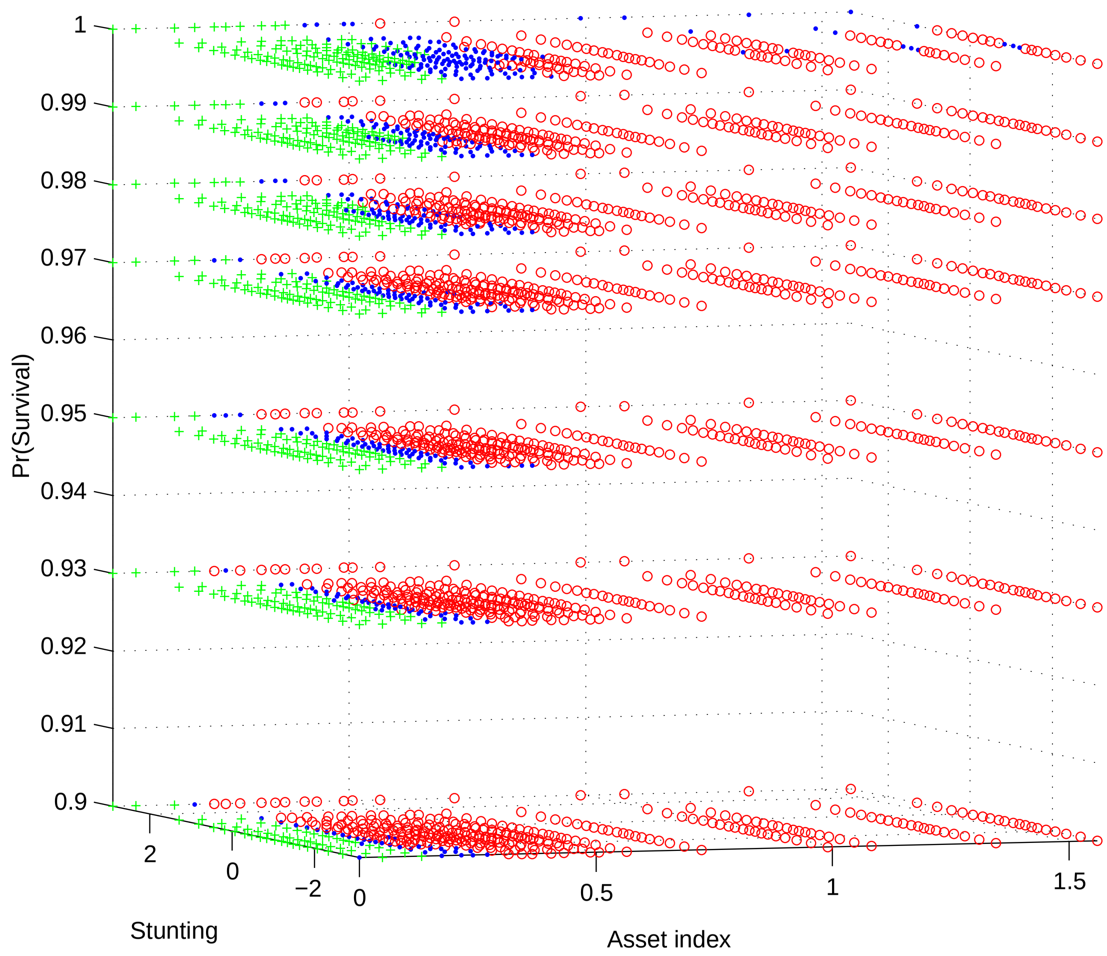

Figure 7: Critical Poverty Frontier, Rural Eastern region *vs.* Urban Northern region in Uganda (critical frontier minus two standard deviations)

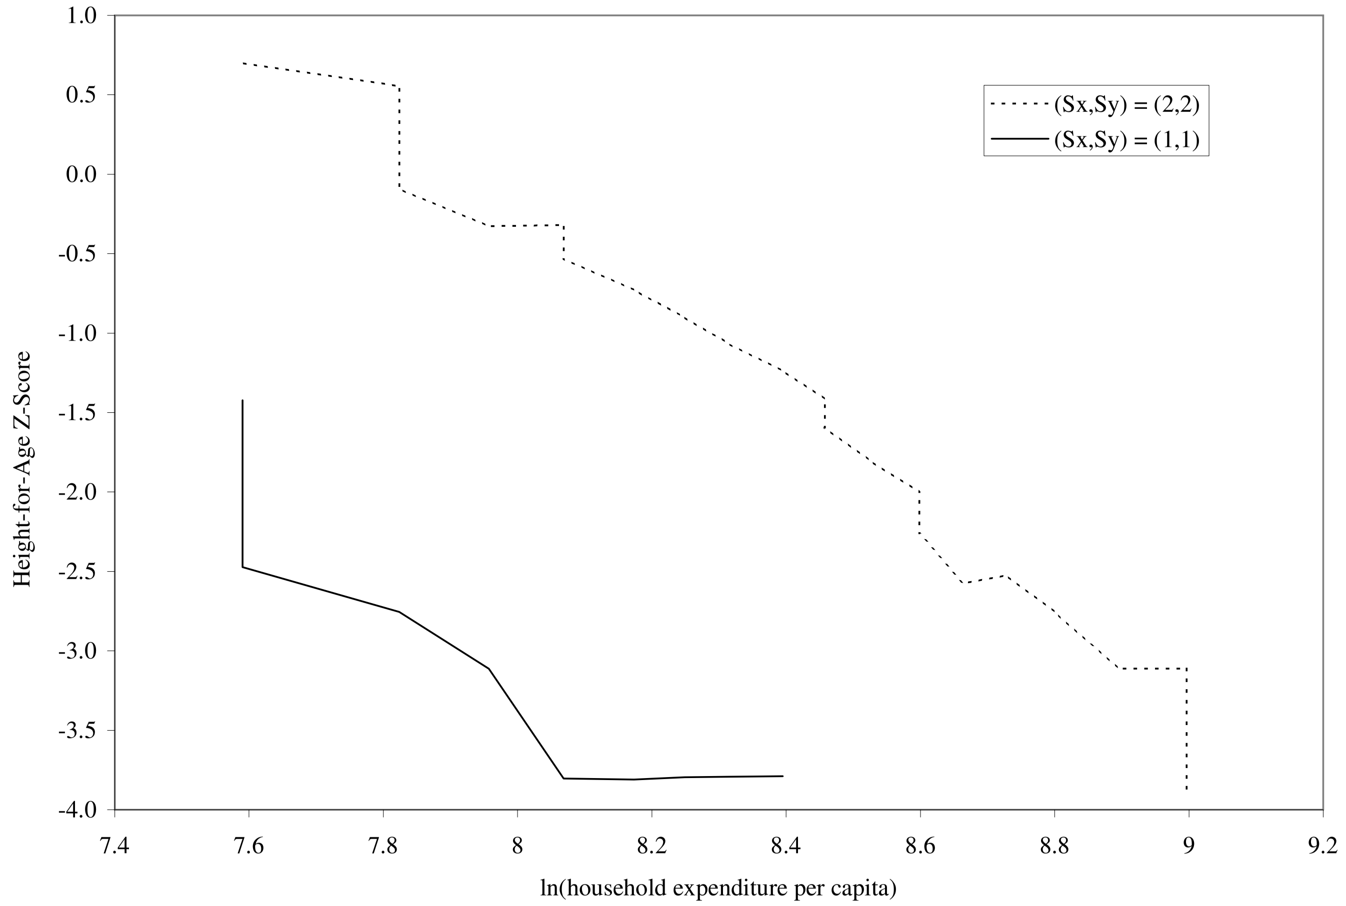

Table 1: Test results for difference between dominance surfaces for urban and rural children in Viet Nam, 1993

log of household expenditure per capita \ height-for-age z-score

	-4.19	-3.64	-3.29	-3.02	-2.82	-2.66	-2.49	-2.31	-2.16	-2.00	-1.83	-1.67	-1.51	-1.32	-1.09	-0.84	-0.49	0.01	0.72	5.47
6.08	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.26	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.37	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.47	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.54	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.61	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.66	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.71	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.77	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.84	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.89	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
6.95	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.03	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.08	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.16	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.27	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.41	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.59	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
7.88	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
9.41	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	0	0

Notes: Sx=1, Sy=1
A negative sign indicates that the urban dominance surface is significantly below the rural one, a positive sign indicates the opposite, and a zero indicates that the difference is not statistically significant.

Table 2: Test results for difference between dominance surfaces for children in Cameroon and Madagascar, 1997

Height-for-age z-score \ Survival probability

	0.83	0.86	0.88	0.89	0.90	0.91		0.99	1.00
-4.19	-	-	0	0	-	-	…	-	-
-3.66	0	0	0	0	-	-	…	-	-
-3.35	0	-	-	-	-	-	…	-	-
-3.13	0	-	-	-	-	-	…	-	-
-2.88	0	-	-	-	-	-	…	-	-
-2.66	-	-	-	-	-	-	…	-	-
-2.50	-	-	-	-	-	-	…	-	-
	…	…	…	…	…	…	…	…	…
0.46	-	-	-	-	-	-	…	-	-
5.39	-	-	-	-	-	-	…	-	.

Notes: 1/ Sx=1, Sy=1

2/ A negative sign indicates that Madagascar's dominance surface is significantly above Cameroon's, a positive sign indicates the opposite, and a zero indicates that the difference is not statistically significant.

3/ The ellipses indicate that all intervening signs are negative.

Table 3: Test results for difference between dominance surfaces for children in Colombia and the Dominican Republic, 1995 and 1996

Height-for-age z-score \ Survival probability	0.906	0.927	0.938	0.947	0.953		0.985	0.987	0.989	0.991	0.995	1.000
-2.85	-	-	-	-	-	…	-	-	-	0	0	0
-2.36	-	-	-	-	-	…	-	-	-	-	0	0
-2.07	-	-	-	-	-	…	-	-	-	-	0	0
-1.85	-	-	-	-	-	…	-	-	-	0	0	0
-1.67	-	-	-	-	-	…	-	-	-	0	0	0
-1.47	-	-	-	-	-	…	-	-	-	0	0	+
-1.33	-	-	-	-	-	…	-	-	-	0	0	+
-1.17	-	-	-	-	-	…	-	-	-	0	+	+
-1.04	-	-	-	-	-	…	-	-	0	0	+	+
-0.92	-	-	-	-	-	…	-	-	-	0	0	+
-0.76	-	-	-	-	-	…	-	-	-	0	+	+
-0.62	-	-	-	-	-	…	-	-	-	0	+	+
-0.49	-	-	-	-	-	…	-	-	-	0	+	+
-0.35	-	-	-	-	-	…	-	-	-	-	+	+
-0.12	-	-	-	-	-	…	-	-	-	-	+	+
0.07	-	-	-	-	-	…	-	-	-	-	0	+
0.34	-	-	-	-	-	…	-	-	-	-	0	+
0.68	-	-	-	-	-	…	-	-	-	-	0	+
1.05	-	-	-	-	-	…	-	-	-	-	-	+
5.92	-	-	-	-	-	…	-	-	-	-	-	0

Notes: Sx=1, Sy=1

A negative sign indicates that the Domincan Republic's dominance surface is significantly above Colombia's, a positive sign indicates the opposite, and a zero indicates that the difference is not statistically significant.
The ellipses indicate that all intervening signs are negative.