Multidimensional Poverty Dominance: Statistical Inference and an Application to West Africa*

Yélé Maweki Batana[†] and Jean-Yves Duclos[‡]

September 24, 2008

Abstract

This paper tests for robust multidimensional poverty comparisons across six countries of the West African Economic and Monetary Union (WAEMU). Two dimensions are considered, nutritional status and assets. The estimation of the asset index is based on two factor analysis methods. The first method uses Multiple Correspondence Analysis; the second is based on the maximization of a likelihood function and on bayesian analysis. Using Demographic and Health Surveys (DHS), pivotal bootstrap tests lead to statistically significant dominance relationships between 12 of the 15 possible pairs of the six WAEMU countries. Multidimensional poverty is also inferred to be more prevalent in rural than in urban areas. These results tend to support those derived from more restrictive unidimensional dominance tests.

Key words: Stochastic dominance, factor analysis, bayesian analysis, multidimensional poverty, empirical likelihood function, bootstrap tests.
JEL Classification: C10,C11,C12,C30,C39,I32.

* We are grateful to Canada's SSHRC, to Québec's FQRSC and to the *Programme canadien de bourses de la francophonie* for financial support. This work was also carried out with support from the Poverty and Economic Policy (PEP) Research Network, which is financed by the Government of Canada through the International Development Research Centre (IDRC) and the Canadian International Development Agency (CIDA), and by the Australian Agency for International Development (AusAID).

† Centre de Recherche Léa-Roback, Direction de la santé publique, Montréal (Québec), H2L 1M3, and Université de Montréal; Tel (514) 528-2400; Fax (514) 528 2453; Email: yele.maweki.batana@umontreal.ca

‡ Département d'économique and CIRPÉE, Pavillon de Sève, Université Laval, Québec, Canada, G1K 7P4; Email: jyves@ecn.ulaval.ca

1 Introduction

The literature on poverty measurement generally follows two approaches. The first one is based on monetary indicators (*e.g.*, Chen and Ravallion 2001 and Atkinson 1998) and essentially treats income or consumption as a unidimensional proxy for welfare; the second approach makes use of a broader set of multidimensional variables (see for instance Streeten, Burki, UL HAQ, Hicks, and Stewart 1981 or Maasoumi 1999). The second approach has gained significantly in popularity since the seminal work of Sen (1985). It also underlies the promotion of the Millenium Development Goals (MDG) by the United Nations, since the MDG focus on deprivation in multiple dimensions.

Although it is indeed now common to assert that poverty is a multidimensional phenomenon, there exist, however, significant difficulties in implementing a truly multidimensional analysis of poverty. In performing such a task, a number of intrinsically arbitrary measurement assumptions are often made, consisting *inter alia* in choosing aggregation procedures across dimensions of well-being, aggregation procedures across across individuals, and multidimensional poverty lines to separate the poor from the non-poor. Each of these choices raises concerns over the possible non-robustness of the results that are obtained.

For instance, an important branch of the literature that considers multiple dimensions of welfare — of which the best-known example is the Human Development Index of the United Nations Development Program (1990) — aggregates simple summary measures of welfare (in terms of life expectancy, literacy, and GDP) into a single one-dimensional index. The across-dimension and across-individual aggregation procedures used in that exercise can easily be criticized — see for instance Kelley (1991) — and several alternative procedures can be (and have been) proposed that lead to alternative views of poverty across time and space.

To allay such concerns over issues of arbitrariness and non-robustness, an alternative to comparing summary indices of multidimensional poverty is to seek poverty comparisons that are valid for a broad class of measurement assumptions. Dominance (or robustness) tests have been in existence in the context of unidimensional comparisons of poverty for many years now (*e.g.*, Atkinson 1987, Foster and Shorrocks 1988a, Foster and Shorrocks 1988b, Anderson 1996, Davidson and Duclos 2000 or Barrett and Donald 2003). It is well-known that one important advantage of such tests is that they are capable of generating poverty comparisons that are robust to the choice of both poverty indices and unidimensional poverty lines. Multidimensional poverty dominance tests have been the object of more recent attention (see for instance Bourguignon and Chakravarty 2002, Atkinson 2003 and Duclos, Sahn, and Younger 2006), even though, as Anderson (2005) points out, adding dimensions to welfare analysis can change one's comparisons of poverty across time and space. In performing multidimensional dominance tests, one is seeking robustness over aggregation procedures across dimensions of welfare, robustness over aggregation procedures across individuals, and robustness over choices of multidimensional poverty lines.

One difficulty with tests for multidimensional poverty dominance is due to the "curse of dimensionality" (see Bellman 1961), a curse that affects all non-parametric comparisons of distributions with multiple variates. This curse is likely to strike in many practical applications of the above multidimensional dominance methodology. The monitoring of the MDG suggests for instance that one should look jointly at a variety of income, health, mortality, educational and environmental indicators. The typical "poverty reduction strategies" drawn by many developing countries draw attention to several dozens of welfare indicators. Comparing the joint distributions of these various indicators across time and space will often prove to be statistically too demanding.

As an alternative to the above, this paper proposes and implements a procedure that stands as a compromise between a desire for greater robustness to measurement assumptions than is usually found in the multidimensional poverty literature, and a practical need for statistical and empirical tractability. To do this, the paper estimates asset indices that incorporate various attributes of difficult-to-aggregate individual indicators of living standards. This is done in the spirit of Sahn and Stifel (2000) and Sahn and Stifel (2003), for instance, which apply factor analysis on multiple welfare indicators to derive a unidimensional welfare index, and on the basis of which they then perform tests for unidimensional dominance in poverty over time and across countries. An important advantage of this approach is to avoid having to use difficult-to-assign index values to various goods and services. The paper subsequently performs a two-dimensional dominance analysis that takes into account an additional indicator of welfare, this time in the dimension of health and nutrition, an indicator that may be distributed quite differently from the above asset index. Hence, we compare welfare using two dimensions, an asset index and health, allowing significant flexibility in mixing up the two dimensions while enforcing sufficient statistical and informational manageability.

Testing for multidimensional dominance then involves comparing joint distribution functions over an infinite number of combinations of possible poverty thresholds in each dimension. This raises obvious computational and statistical difficulties. Another contribution of the paper is thus to extend a recently proposed statistical inference procedure for univariate distributions to the case of multivariate distributions. This is done by applying the empirical likelihood method suggested by Davidson and Duclos (2006) to test the existence of two-dimensional poverty dominance relationships. The outcome is the derivation of an intersection-union test procedure that enables inferring strict dominance relationships in a multivariate context.

The above measurement and statistical methodology is then applied to comparing poverty across six members of the West African Economic and Monetary Union: Benin, Burkina Faso, Côte d'Ivoire, Mali, Niger, and Togo. Wide disparities are observable across these countries, with Côte d'Ivoire accounting for nearly 40% of the output of the Union. Poverty is also compared within each country across rural and urban areas, in part to check whether the usual monetary comparisons that show lower poverty in urban areas also carry over to the case of multiple dimensions. It is indeed possible that, even with low overall average incomes, rural inhabitants may possess greater assets, may have better and more direct access to agricultural produce, and may be better nourished than city dwellers, who must frequently pay a higher price for what may be lower-quality foodstuffs.

The rest of the paper is as follows. Section 2 presents the methodology for estimating the asset index. Section 3 explains what two-dimensional stochastic dominance is and sets up the techniques for testing it. The data description and the empirical results are found in Section 4. Section 5 concludes.

2 Estimating the welfare indicators

We first define the two indicators of welfare used for the paper's multidimensional poverty analysis. The first indicator is a measure of nutrition and health; the second measures assets and is derived from two estimation procedures, *i.e.*, an inertia method based on Multiple Correspondence Analysis (MCA) and a factor analysis procedure using likelihood.

2.1 Calculating an indicator of health and nutrition

Several measures of health and nutrition are used in the literature, the main ones being weight-for-height indices, height-for-age indices, and weight-for-age indices. These indices are obtained from comparisons of weights or heights with the mean or median value of a reference population. For children, Sahn and Stifel (2002) argue for the use of a height-for-age index since it is not as much affected as other indices by episodes of stress, diarrhoea, malaria, or other conditions that may temporarily modify the health and nutritional status of an individual. Instead, height-for-age indices tend to capture the cumulative impact on health of longer-term factors, such as average socio-economic conditions and public health policy — vaccination programs, efforts to combat endemic diseases and other chronic illness, or sanitation programs for instance.

A height-for-age Z_score for a child i can be computed as

$$Z\_score_i = \frac{T_i - T_{median}}{\sigma_T}, \tag{1}$$

where T_i is the child's body-height, T_{median} is the median body-height of a healthy and well-nourished child from the reference population used by the United States National Center for Health Statistics, and σ_T is the standard deviation of body-heights in the reference population. The children in question are between 3 and 35 months. By convention, a child with a Z_score falling below -2 (a nutritional poverty "threshold") is usually deemed to be suffering from malnutrition.

2.2 Estimating the asset index

2.2.1 An inertia approach

An inertia approach is used to derive an asset index. Let each of N individuals, indexed $i = 1, \ldots, N$, exhibit J welfare attributes, $j = 1, \ldots, J$. These N individuals can be represented by a cloud of points around a centroid (the weighted means) in the space of the J attributes, with each point having some weight. The total inertia of the cluster of points is the weighted sum of the distance of each point from the centroid.

The main issue is how to proceed to the estimation of an asset index for each household using a weighted sum of the welfare attributes. Let X_i be the asset index for individual i, x_{ij} be his endowment of attribute j, and α_j be the weight assigned to each attribute. X_i is then given by

$$X_i = \alpha_1 x_{i1} + \ldots + \alpha_J x_{iJ}. \tag{2}$$

In order to check for robustness over the choice of the inertia method for the data reduction operation implicit in (2), two methods are used, both being suitable when we deal with qualitative variables. The first one, a MCA procedure, is well-known and will not be described in detail here (see for instance Greenacre 1993 and Greenacre and Blasius 2006). The second method, which involves a maximum likelihood procedure, has proven to be quite useful in social sciences (such as psychology) and is based on confirmatory factor analysis with qualitative variables. Since it is less familiar to economists, we outline it below.

2.2.2 factor analysis approach

The literature has essentially proposed two procedures, one based on an underlying response variable and another one based on a response function (see for instance Moustaki 2000 and Jöreskog and Moustaki 2001). With the first procedure, each of the qualitative indicators is supposed to be generated by an unobserved, continuous response variable (a latent variable) that is distributed normally. For the second procedure, a conditional distribution for each possible configuration of J-dimensional responses is specified as a function of the latent factors, assuming that the responses to the various variables are independent — see Bartholomew (1983), Bartholomew (1984) and Moustaki (2000), among others. Jöreskog and Moustaki (2001) suggest three estimation methods for each of these two procedures, each with the benefit of estimating the parameters in a single step. We choose the so-called Underlying Multivariate Normal (UMN) method.

To see this better, consider the following model,

$$x_{ij}^* = \lambda_j f_i + \varepsilon_{ij}, \quad j = 1, 2, \ldots, J \quad \text{and} \quad i = 1, 2, \ldots, N. \tag{3}$$

The factor f_i, which is specific to each individual i, captures the individual's unobserved welfare level. The error term, ε_{ij}, also unobserved, is specific to each household and to each variable. The difference with the linear model of Sahn and Stifel (2000) and Sahn and Stifel (2003) is that the continuous response variable x_{ij}^* is unobserved in the present case. The model in (3) also differs from standard limited dependent variable models in that f_i is unobserved. Dropping the index i (without loss of generality), x_j^* and x_j (which is observed) are linked as follows:

$$x_j = a \Longleftrightarrow \gamma_{a-1}^{(j)} < x_j^* \leq \gamma_a^{(j)}, \quad a = 1, 2, \ldots, m_j, \tag{4}$$

where $\gamma_0^{(j)} = -\infty$, $\gamma_1^{(j)} < \gamma_2^{(j)} < \ldots < \gamma_{m_j-1}^{(j)}$, $\gamma_{m_j}^{(j)} = +\infty$ are the threshold parameters. Thus, for m_j categories of response for a given variable, there are $m_j - 1$ threshold parameters. Since the mean and the variance of x_j^* are not identified, we set them equal to 0 and 1, respectively. We also assume that f and ε_j are independently and normally distributed with $f \sim N(0, 1)$ and $\varepsilon_j \sim N(0, \psi_j)$. Since the first two moments of x_j^* have been normalized to 0 and 1 respectively, we have that $\psi_j = 1 - \lambda_j^2$. The distribution of $x_1^*, \ldots, x_J^*$ is then multivariate normal with mean zero, variance one, and a correlation matrix $\Gamma = (\rho_{jk})$, where $\rho_{jk} = \lambda_j \lambda_k$. The parameters to be estimated are the threshold parameters $\gamma_a^{(j)}$ and the coefficients of the factors λ_j, with $j = 1, 2, \ldots, J$ and $a = 1, 2, \ldots, m_j - 1$. The total number of parameters to be estimated is then $\Sigma_{j=1}^J m_j$.

Let θ be the set of parameters to be estimated and r be one of the possible configurations of J respective responses on the level of the J qualitative variables. The probability of the realization of that specific configuration is then given by:

$$\begin{aligned} \pi_r(\theta) &= \Pr(x_1 = a_1, x_2 = a_2, ..., x_J = a_J), \\ &= \int_{\gamma_{a_1-1}^{(1)}}^{\gamma_{a_1}^{(1)}} \int_{\gamma_{a_2-1}^{(2)}}^{\gamma_{a_2}^{(2)}} \cdots \int_{\gamma_{a_J-1}^{(J)}}^{\gamma_{a_J}^{(J)}} \phi_J\left(\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_J | \Gamma\right) d\varepsilon_1 d\varepsilon_2 \ldots d\varepsilon_J. \end{aligned} \tag{5}$$

ϕ_J is the normal multivariate density function with J dimensions. The log-likelihood function is then given by:

$$L_{UMN}(\theta) = \sum_r P_r \ln \pi_r(\theta), \tag{6}$$

where $P_r = \frac{n_r}{N}$, n_r is the number of realizations of configuration r observed in the sample, and N is sample size. Maximizing $L(\theta)$ with respect to θ, we obtain a set of full information maximum likelihood estimators.

When $J > 4$, however, the numerical maximization of (6) can be computationally demanding. This is why Jöreskog and Moustaki (2001) advocate estimation based on an Underlying Bivariate Normal (UBN) procedure. Rather than calculating the probabilities of all of the various configurations of the J responses, one is interested in $\pi_a^{(j)}$ — the standard probability of obtaining a response in category a for variable j — and in the probabilities of simultaneously having a response in category a for variable j and a response in category b for variable h. These expressions are respectively given by:

$$\begin{aligned} \pi_a^{(j)}(\theta) &= \int_{\gamma_{a-1}^{(j)}}^{\gamma_a^{(j)}} \phi_1\left(u\right) du, & (7) \\ \pi_{ab}^{(jh)}(\theta) &= \int_{\gamma_{a-1}^{(j)}}^{\gamma_a^{(j)}} \int_{\gamma_{b-1}^{(h)}}^{\gamma_b^{(h)}} \phi_2\left(u, v | \rho_{jh}\right) du\, dv, & (8) \end{aligned}$$

where $\phi_1\left(u\right)$ is a standard normal density function and $\phi_2\left(u, v|\rho\right)$ is a normal bivariate density function with correlation coefficient ρ. The parameters are then estimated by maximizing:

$$L_{UBN}(\theta) = \sum_{j=1}^{J} \sum_{a=1}^{m_j} P_a^{(j)} \ln \pi_a^{(j)}(\theta) + \sum_{j=2}^{J} \sum_{h=1}^{j-1} \sum_{a=1}^{m_j} \sum_{b=1}^{m_h} P_{ab}^{(jh)} \ln \pi_{ab}^{(jh)}(\theta). \tag{9}$$

Even though this technique has no clear theoretical basis, Jöreskog and Moustaki (2001) find that it yields the same results as full information techniques based on the entire response function method. The technique is, however, considerably less demanding in terms of processing time and it works well even when the number of qualitative variables is large.

The next step involves estimating the factor scores f. Note that in the case of ordered variables, there is no linear relationship between the factors and the observed variables. To estimate the score, Shi and Lee (1997) propose a Bayesian approach based on the *a posteriori* distribution of factors. In the following, f is a scalar, λ, x and x^* are $(J \times 1)$ vectors, and ψ is a $(J \times J)$ diagonal matrix. If we let $p(f)$ be the density function of f, $p(x^*|f)$ be the density function of x^* conditional on f, and $\Pr(x|f)$ be the conditional probability of x given f, then, according to Bayes' theorem, the conditional distribution of f given x is the solution to:

$$p(f|x) = \frac{p(f)\Pr(x|f)}{\int_R p(f)\Pr(x|f)df}. \tag{10}$$

Using (10), we can estimate a score f for each individual i as

$$X_i = E(f|x) = \frac{1}{B}\lambda'\psi^{-1}x_W^*. \tag{11}$$

The expressions for B and x_W^*, and the steps needed to obtain them, are shown in the appendix.

3 Multidimensional stochastic dominance

The dominance techniques used here are drawn from Duclos, Sahn, and Younger (2006) and are multidimensional extensions to the stochastic dominance techniques developed by Atkinson (1987), Foster and Shorrocks (1988a) and Foster and Shorrocks (1988b) for a one-dimensional framework. They make ordinal poverty comparisons possible over classes of procedures for aggregating across dimensions and across individuals. They also allow for robustness over areas of possible multidimensional poverty "frontiers" — analogous to the usual unidimensional poverty lines. We outline these techniques briefly below.

3.1 First-order dominance

Duclos, Sahn, and Younger (2006) start by defining a generic additive multidimensional poverty index as

$$P(\lambda) = \int\int_{\Lambda(\lambda)} \pi(x_1, x_2; \lambda)\, dF(x_1, x_2), \tag{12}$$

where $\lambda(x_1, x_2)$ is a function that captures overall welfare (analogous to a utility function), $\lambda(x_1, x_2) = 0$ is a poverty "frontier" that separates the rich from the poor, and $\Lambda(\lambda)$ is the (x_1, x_2) area defined as $\lambda(x_1, x_2) \le 0$ within which the set of poor people can be found.

Note that the definition of $\lambda\left(x_1, x_2\right)$ is general enough to encompass what are known as union, intersection or intermediate definitions of the poor. This is illustrated in Figure 1, where x_1 and x_2 are two dimensions of welfare. $\lambda_1(x_1, x_2)$ provides an "intersection" definition of poverty: it considers someone to be in poverty only if he is poor in *both* of the two dimensions. $\lambda_2(x_1, x_2)$ gives a union poverty index: it considers someone to be in poverty if he is poor in *either* of the two dimensions. $\lambda_3(x_1, x_2)$ provides an intermediate definition: someone can be poor even if $x_1 > Z_1$, if his x_2 value is sufficiently low to lie to the left of $\lambda_3(x_1, x_2) = 0$. Alternatively, someone can be non-poor even if $x_1 < Z_1$ if his x_2 value is sufficiently high to lie to the right of $\lambda_3(x_1, x_2) = 0$.

Figure 1: Definitions of union, intersection, and intermediate poverty

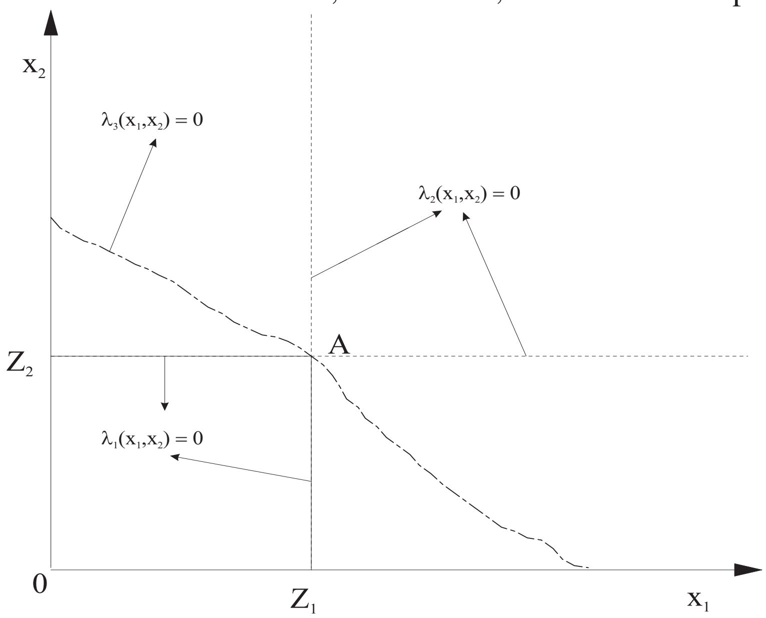

The dominance test that Duclos, Sahn, and Younger (2006) propose then uses a two-dimensional extension of the well-known FGT index (Foster, Greer, and Thorbecke 1984)

$$P_{\alpha_1,\alpha_2}\left(z_1, z_2\right) = \int_0^{z_1}\int_0^{z_2}\left(z_1 - x_1\right)^{\alpha_1}\left(z_2 - x_2\right)^{\alpha_2} dF\left(x_1, x_2\right), \tag{13}$$

where F is the bivariate cumulative distribution function and α_1 and α_2 are non-negative parameters that capture aversion to inequality in poverty in each of the two dimensions. With dominance orders s_1 and s_2 set to $s_1 = 1 + \alpha_1$ and $s_2 = 1 + \alpha_2$, plotting (13) over an area of z_1 and z_2 provides a dominance surface for a distribution F. The difference in that surface between distributions F and G is then given by

$$\Delta P^{s_1 s_2}(z_1, z_2) = \int_0^{z_1} \int_0^{z_2} (z_1 - x_1)^{s_1 - 1} (z_2 - x_2)^{s_2 - 1} \, d(F - G)(x_1, x_2). \quad (14)$$

To show how (14) can serve to order distributions in terms of multidimensional poverty, Duclos, Sahn, and Younger (2006) uses (12) to define the following first-order class of poverty indices $\Psi^{1,1}(\lambda^*)$:

$$\Psi^{1,1}(\lambda^*) = \left\{ P(\lambda) \;\middle|\; \begin{array}{c} \Lambda(\lambda) \subset \Lambda(\lambda^*) \\ \pi(x_1, x_2; \lambda) = 0 \text{ if } \lambda(x_1, x_2) = 0 \\ \frac{\partial \pi(x_1, x_2; \lambda)}{\partial x_1} \leq 0 \text{ and } \frac{\partial \pi(x_1, x_2; \lambda)}{\partial x_2} \leq 0 \; \forall \, x_1, x_2 \\ \frac{\partial^2 \pi(x_1, x_2; \lambda)}{\partial x_1 \partial x_2} \geq 0 \; \forall \, x_1, x_2 \end{array} \right\} \quad (15)$$

The first row of (15) defines the maximum set of poor people. The second row assumes continuity of the poverty indices along the poverty frontier. The third row follows from an axiom of *monotonicity* and states that the indices should be weakly decreasing in the attributes x_1 and x_2. The last row reflects an axiom of *attribute substitutability* — essentially saying that the greater the value of an attribute, the lesser the impact on poverty of an increase in the value of the other attribute. Theorem 1 in Duclos, Sahn, and Younger (2006) on first-order dominance then says that all of the multidimensional poverty indices in the class of measures $\Psi^{1,1}(\lambda^*)$ will be greater in F than in G if and only $\Delta P^{1,1}(z_1, z_2) > 0 \, \forall \, (z_1, z_2) \in \Lambda(\lambda^*)$.

3.2 Higher-order stochastic dominance

It is possible to derive higher-order dominance conditions, but this requires further assumptions on the sign of the derivatives of order higher than in (15). The order of dominance can be increased in either of the dimensions individually, or in both simultaneously, leading for example to such classes as $\Psi^{2,1}(\lambda^*)$, $\Psi^{1,2}(\lambda^*)$ or $\Psi^{2,2}(\lambda^*)$. $\Psi^{2,1}(\lambda^*)$ is for instance defined as

$$\Psi^{2,1}(\lambda^*) = \left\{ P(\lambda) \;\middle|\; \begin{array}{c} P(\lambda) \in \Psi^{1,1}(\lambda^*) \\ \frac{\partial \pi(x_1, x_2; \lambda)}{\partial x_1} = 0 \text{ if } \lambda(x_1, x_2) = 0 \\ \frac{\partial^2 \pi(x_1, x_2; \lambda)}{(\partial x_1)^2} \geq 0 \; \forall \, x_1 \\ \frac{\partial^3 \pi(x_1, x_2; \lambda)}{(\partial x_1)^2 \partial x_2} \leq 0 \; \forall \, x_1, x_2 \end{array} \right\} \quad (16)$$

The first row imposes compliance with the conditions of belonging to the class $\Psi^{1,1}(\lambda^*)$. The second row says that the first derivative with respect to x_1 should be continuous along the poverty frontier. The third imposes the well-known Pigou-Dalton *principle of transfer* on attribute x_1: it says that the poverty impact of increasing x_1 should decrease with x_1, or alternatively that an equalizing transfer in the x_1 dimension should diminish poverty. The last row assumes that the equalizing effect of such a transfer should declines with x_2; said differently, the greater the value of x_2, the lesser the importance of inequality in the dimension of x_1.

Theorem 2 of Duclos, Sahn, and Younger (2006) then says that G dominates F in poverty over for the class $\Psi^{2,1}(\lambda^*)$ of poverty indices — namely, all of the multidimensional poverty indices in the class of measures $\Psi^{2,1}(\lambda^*)$ will be greater in F than in G — if and only if $\Delta P^{2,1}(z_1, z_2) > 0 \; \forall \; (z_1, z_2) \in \Lambda(\lambda^*)$.

3.3 Statistical inference

From the above, non-dominance of distribution F by distribution G implies that there exists a point (z_1, z_2) in $\Lambda(\lambda^*)$ for which $\Delta P^{s_1 s_2}(z_1, z_2) \leq 0$. This suggests the following set of null and alternative hypotheses for tests of multidimensional dominance:

$$
\begin{aligned}
H_0 &: \Delta P^{s_1 s_2}(z_1, z_2) \leq 0 \text{ for some } (z_1, z_2) \text{ in } \Lambda(\lambda^*) \\
\text{versus} & \\
H_1 &: \Delta P^{s_1 s_2}(z_1, z_2) > 0 \text{ for all } (z_1, z_2) \text{ in } \Lambda(\lambda^*) \; . \qquad (17)
\end{aligned}
$$

H_0 is a null of non-dominance of F by G. If this null is rejected, then all that is logically left is dominance of F by G, which is the alternative H_1.

To test for H_0 against H_1, we extend the use of the empirical likelihood ratio (ELR) statistic suggested by Davidson and Duclos (2006) for univariate distributions to the case of multivariate distributions. The ELR statistic captures the "distance" between the empirical distributions and the null hypothesis of non-dominance. It equals the difference between the unconstrained empirical likelihood of the distributions and the empirical likelihood constrained by H_0.

To see this more clearly, let n_i^F represent the number of sample observations that equal $\left(x_{i1}^F, x_{i2}^F\right)$, where $\left(x_{i1}^F, x_{i2}^F\right)$ are the values of indicators 1 and 2 taken by the i^{th} observation of a sample of N_F independently and identically distributed observations drawn from distribution F — and analogously for n_j^G. Also, let $I(\cdot)$ be an indicator function assuming the value one when the argument is true, and zero otherwise, and let P_i^F and P_j^G be the empirical probabilities for the observations of samples from F and G respectively. The ELR statistic (in log form) is then obtained by maximizing the following empirical likelihood function

$$\max_{P_i^F, P_j^G} \sum_i n_i^F \log P_i^F + \sum_j n_j^G \log P_j^G \tag{18}$$

subject to

$$\sum_i P_i^F = 1, \ \sum_j P_j^G = 1 \tag{19}$$

and constrained — or not — by

$$\begin{aligned} & \sum_i P_i^F (z_1 - x_{i1}^F)^{s_1 - 1} (z_2 - x_{i2}^F)^{s_2 - 1} I \left(x_{i1}^F \leq z_1, x_{i2}^F \leq z_2 \right) \\ \leq & \sum_j P_j^G (z_1 - x_{j1}^G)^{s_1 - 1} (z_2 - x_{j2}^G)^{s_2 - 1} I \left(x_{j1}^G \leq z_1, x_{j2}^G \leq z_2 \right) \end{aligned} \tag{20}$$

for some (z_1, z_2) in $\Lambda(\lambda^*)$.

It can be checked that the maximization of (18) unconstrained by (20) is given by $N \log N - N_F \log N_F - N_G \log N_G$. As for the constrained maximum, if F dominates G in the sample, then there is no cost to imposing (20); in this case, the constraint is not binding and the ELR statistic equals zero. We cannot then reject H_0.

If instead G dominates F in the sample, we may distinguish between first-order and higher-order dominance. If $s_1 = s_2 = 1$, the constraint in (20) becomes:

$$\sum_i P_i^F I \left(x_{i1}^F \leq z_1, x_{i2}^F \leq z_2 \right) = \sum_j P_j^G I \left(x_{j1}^G \leq z_1, x_{j2}^G \leq z_2 \right). \tag{21}$$

Proceeding to (18) subject to (21) yields the following empirical probabilities:

$$\begin{aligned} P_i^F &= \frac{n_i^F I_i(z_1, z_2)}{\varpi} + \frac{n_i^F (1 - I_i(z_1, z_2))}{\psi}, \\ P_j^G &= \frac{n_j^G I_j(z_1, z_2)}{N - \varpi} + \frac{n_j^G (1 - I_j(z_1, z_2))}{N - \psi}, \end{aligned} \tag{22}$$

with

$$I_i(z_1,z_2) = I\left(x_{i1}^F \le z_1, x_{i2}^F \le z_2\right), I_j(z_1,z_2) = I\left(x_{j1}^G \le z_1, x_{j2}^G \le z_2\right), \quad (23)$$

$$\varpi = \frac{N \times N_F(z_1,z_2)}{N_F(z_1,z_2) + N_G(z_1,z_2)}, \psi = \frac{N \times M_F(z_1,z_2)}{M_F(z_1,z_2) + M_G(z_1,z_2)}, \quad (24)$$

$$N_F(z_1,z_2) = \sum_i n_i^F I_i(z_1,z_2), N_G(z_1,z_2) = \sum_j n_j^F I_j(z_1,z_2), \quad (25)$$

$$N = N_F + N_G, \quad (26)$$

$$M_F(z_1,z_2) = N_F - N_F(z_1,z_2), M_G(z_1,z_2) = N_G - N_G(z_1,z_2). \quad (27)$$

The ELR statistic is then obtained by the difference between the unconstrained maximum and the maximum over $(z_1,z_2) \in \Lambda(\lambda^*)$ of the constrained likelihood (given by $\sum_i n_i^F \log P_i^F + \sum_j n_j^G \log P_j^G$ using (22)). This yields:

$$\frac{1}{2}LR(z_1,z_2) = \left\{ \begin{array}{c} N\log N - N_F \log N_F - N_G \log N_G \\ +N_F(z_1,z_2)\log N_F(z_1,z_2) + N_G(z_1,z_2)\log N_G(z_1,z_2) \\ +M_F(z_1,z_2)\log M_F(z_1,z_2) + M_G(z_1,z_2)\log M_G(z_1,z_2) \\ -(N_F(z_1,z_2)+N_G(z_1,z_2))\log(N_F(z_1,z_2)+N_G(z_1,z_2)) \\ -(M_F(z_1,z_2)+M_G(z_1,z_2))\log(M_F(z_1,z_2)+M_G(z_1,z_2)). \end{array} \right\} \quad (28)$$

It is then a matter of notation to use Theorem 1 of Davidson and Duclos (2006) to show that (28) is asymptotically equivalent to the square of an asymptotically normally distributed t statistic used by Kaur, Rao, and Singh (1994) and given in our context by

$$\frac{N_F N_G \left(\widehat{F}(z_1,z_2) - \widehat{G}(z_1,z_2)\right)^2}{N_G \widehat{F}(z_1,z_2)\left(1-\widehat{F}(z_1,z_2)\right) + N_F \widehat{G}(z_1,z_2)\left(1-\widehat{G}(z_1,z_2)\right)}, \quad (29)$$

where $\widehat{F}(z_1,z_2)$ and $\widehat{F}(z_1,z_2)$ are the empirical distribution functions of F and G, respectively. We can also show that, on the frontier of the null H_0 of non-dominance, both the ELR statistic and (29) are asymptotically pivotal, that is, they follow the same asymptotic distribution for all configurations of the population distributions that lie on the frontier.

As in Davidson and Duclos (2006), therefore, we can perform bootstrap tests to yield more satisfactory inference than tests based solely on the asymptotic distribution of the ELR statistic. To do this (again, this needs to be done only when dominance of F by G exists in the sample), we compute the maximum of the ELR statistic (28) over $(z_1,z_2) \in \Lambda(\lambda^*)$ — denote this maximum as $LR_0(z_1,z_2)$

— and calculate the associated probabilities in (22). These probabilities are then used to generate a certain number (399 in the illustration below) of bootstrap samples for both distributions. For each pair of such bootstrap samples, a new ELR statistic in (28) is computed. A bootstrap value p is computed as the proportion of bootstrap ELR statistics that exceed $LR_0\left(z_1, z_2\right)$.

For higher-order dominance, the same procedure can be followed, except that there exists no analytical solution analogous to (22). Details of the computations are provided in the Appendix for such cases.

4 Empirical comparisons of welfare in West Africa

4.1 The data

We apply the above methodology to comparing multidimensional poverty across countries that are members of the West African Economic and Monetary Union (WAEMU) and for period of the mid-90s (1996–98). The data come from nationally representative Demographic and Health Surveys (DHS) of urban and rural households, and cover six countries from West Africa: Benin, Burkina Faso, Côte d'Ivoire, Mali, Niger, and Togo. Senegal, an important member of WAEMU, was excluded because nutritional data were not collected for the period under consideration. Basic descriptive information on the the surveys used can be found in Table 1. All estimates take into account the sampling weight of each observation.

The DHS surveys provide the necessary information required for calculating a nutritional index (Z_score) and an asset index (X). To compute the Z_score, we use data on the body-height, age, and sex of children, as well as standard values for the reference children population. We only count women for whom nutritional data was collected on one child in the household. The asset index then takes into account household well-being, and the nutritional index is based on child well-being.

As to X, we estimate it from information on ownership of durable goods (radio, television, refrigerator, bicycle, motorcycle, car) and on access to other goods and services (electricity, type of toilette, quality of flooring, potable water, education). All variables are qualitative, which is why MCA and the UBN procedures are used. We use all of the six samples combined to generate the factor scores. For the UBN method, we first estimate the parameters by maximizing (9). These estimates then contribute to the calculations of the factor scores in (11). For the Monte Carlo simulation, 100,000 random vectors are generated from a uniform distribution.

Descriptive statistics on these indices are presented in Tables 2 and 3. On average, Côte d'Ivoire, Togo, and Mali post the best welfare levels, while Burkina

Faso and Niger are weakest.

Table 4 presents a sensitivity analysis on the two methods used to estimate the asset index. Whether the estimated indices describe welfare is informed by whether the percentage of households not affording a good or a service declines with movement from a lower to a higher quartile of the asset indices. Both the MCA and the UBN procedures seem to work reasonably well. Table 5 also demonstrates that these two indices are highly (though not perfectly) correlated, with a correlation coefficient of 96 per cent for the entire sample. The UBN estimator is henceforth used for presenting the results of the dominance analysis, although dominance is not said to be accepted below unless it is confirmed by both the UBN and the MCA asset indices.

4.2 Results of the dominance tests

To test for multidimensional poverty dominance, one should in theory test over the full set of points in $\Lambda(\lambda^*)$. This, however, proves computationally tedious when sample sizes are large. We therefore use a grid of points (z_1, z_2), rather than consider all points in the two distributions F and G, with z_1, $\left[z_1^-, z_1^+\right]$, and z_2, $\left[z_2^-, z_2^+\right]$. To constitute the grid, we consider 20 quantiles for the asset index and 10 quantiles (deciles) for the nutritional index, which we determine after merging the two distributions to be compared. Moreover, rather than consider the quantiles as such, the grid is created by taking the mean of each quantile. This yields a total of 200 points for each comparison. All of the checks we have made suggest, however, that the results below are robust to increasing the number of points with that grid.

Table 6 presents the country-wise results of the first-order dominance tests. The first country represents distribution F, and the second distribution G. The test is conducted against the null hypothesis that G does not dominate F. With the six countries, there are 15 possible comparisons, and thus 15 possible dominance relationships. The results in Table 6 reveal the existence of 12 statistically significant dominance relationships with p values lower than 10%, 10 with p values lower than 5%, and 8 with p values lower than 1%.

Aside from the case Burkina-CI (CI stands for Côte d'Ivoire), which is a dominance relationship that extends over the full grid of $\Lambda(\lambda^*)$ described above, the 11 other dominance relationships are more limited, in the sense that it is necessary to exclude certain points on the grid (in the lower or upper extremities) to obtain dominance. Still, dominance in these other relationships is both statistically quite strong and normatively robust to both the choice of a wide area of possible poverty frontiers (so long as they fall in $\Lambda(\lambda^*)$) and to the choice of multidimensional poverty indices within the class $\Psi^{1,1}(\lambda^*)$).

The three cases of non-dominance shown in Table 6 correspond to situations in which there are several points of intersection between the two distributions, such that it proves impossible to obtain significant differences in dominance surfaces. Côte d'Ivoire dominates all countries, followed by Togo, which dominates 3 countries, to wit Benin, Burkina Faso, and Niger. Benin and Mali dominate Burkina Faso and Niger. We observe no dominance between Mali and Togo, Mali and Benin, and Burkina and Niger.

Table 7 presents similar statistics for comparisons of welfare in rural and urban areas *within* each of the six different countries, and for all countries taken together. It clearly shows that the urban area dominates the rural one, for every country as well as for the global sample. All p values are below 5%, and 4 of the 7 dominance relationships for Burkina Faso, Mali, Togo, and the whole sample show unrestricted dominance, namely, over the the widest possible $\left[z_1^-, z_1^+\right] \otimes \left[z_2^-, z_2^+\right]$ and thus over the entire range of possible poverty frontiers. The other dominance relationships are more restricted but still very robust.

To order the countries that are not ranked by Table 6, we may proceed to tests for higher-order dominance. Table 8 shows the results. The nulls of non-dominance still cannot be rejected for comparisons of Mali and Togo, even at order 3. This is because the curves of the two distributions intersect for each of these orders. This is not the case for the two other relationships, since Mali now dominates Benin in the second order, while Niger dominates Burkina in the third order.

Figure 2: Diagram of dominance among the countries of WAEMU

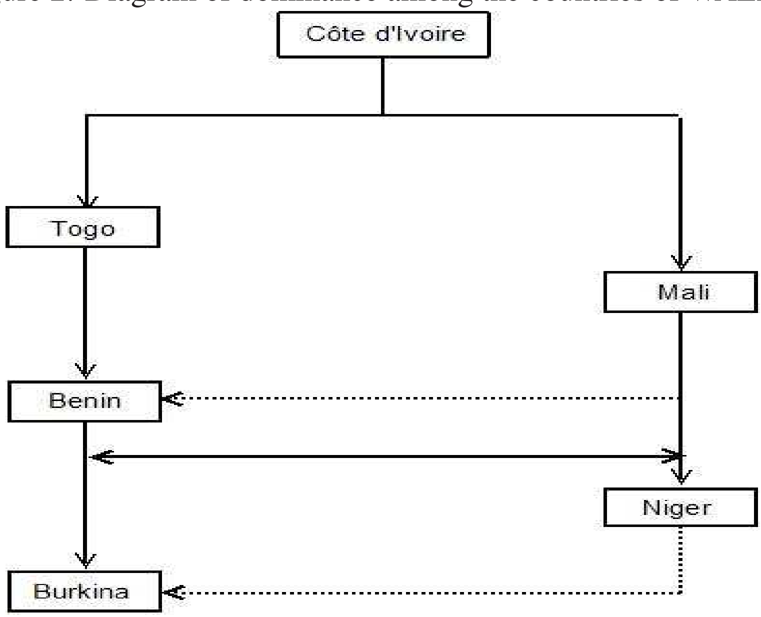

Figure 2 depicts the classification of the countries in terms of dominance. A solid arrow reflects first-order dominance, while a dashed arrow represents higher-order dominance. The position of each country vis-à-vis the peak reveals its position in terms of welfare. While Côte d'Ivoire's presence at the peak reflects the fact that it has the lowest level of multidimensional poverty, Burkina Faso is at the bottom with the highest level.

Figure 3: Diagrams of dominance of WAEMU countries, by area — rural on the left, urban on the right

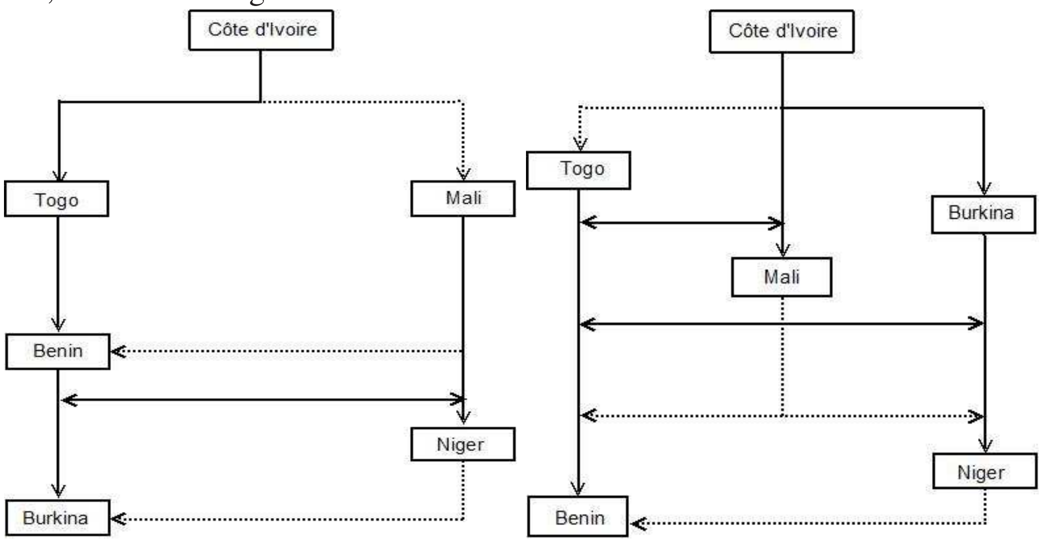

We use this tool to fine-tune our analysis by decomposing dominance with respect to the rural-urban location of individuals. We thus obtain two diagrams in Figure 3, representing country-wise dominance relationships for each of the rural and urban areas, with the left-hand side representing dominance across rural areas, and the right-hand side dominance across urban areas. Note that the dominance relationships across rural areas are practically identical to those in Figure 2. The only difference is that the Côte d'Ivoire's rural zone dominates that of Mali only in the second, as opposed to the first, order.

Urban dominance results in Figure 3 are much more surprising. Burkina, which was previously dominated by all countries, is now dominated only by Côte d'Ivoire, and it dominates Niger. Benin, which significantly dominated Niger and Burkina, is now poorer than all in terms of its urban population. This situation reveals an imbalance between urban and rural standards of living in some countries, notably in Burkina Faso, where urban monetary poverty has regularly been estimated to be far less severe than the rural one.

This also points to the importance of disaggregating poverty comparisons within countries. Since the distribution of welfare across socio-economic groups may differ significantly across countries, cross-country comparisons of national poverty can indeed hide important discrepancies within countries. Uncovering these discrepancies helps understand the context-specific sources of national poverty.

5 Conclusion

Stochastic dominance has almost always been analyzed in the framework of univariate comparisons of welfare. In most cases, formal statistical tests have not been applied to the empirical comparisons. Sub-Saharan Africa, where poverty is in all likelihood the greatest, has relatively rarely been the object of empirical tests for poverty dominance, especially using data that are readily comparable across countries and time.

This paper attempts to move forward in all of these aspects of performing welfare comparisons. Drawing on recent work on making robust comparisons of multidimensional poverty, two dimensions of welfare are considered and compared, nutritional status and assets, using a set of comparable variables drawn from the easily accessible Demographic and Health Surveys of six West-African countries. The estimation of the asset index is based on two factor analysis methods. The first method uses Multiple Correspondence Analysis, and the second one is based on confirmatory factor analysis with qualitative variables.

Statistical inference for the paper's multidimensional poverty comparisons uses a multivariate extension of a recently proposed empirical likelihood ratio test. Because the test statistic that is derived and used is asymptotically pivotal, we are able to perform bootstrap tests that can be expected to yield more satisfactory inference than the usually considered tests that are based solely on analytic asymptotic distributions.

The statistical multidimensional dominance tests we perform across Benin, Burkina Faso, Côte d'Ivoire, Mali, Niger, and Togo confirm the usual (unidimensional) result that poverty is more pronounced in the countryside than in the cities. They also lead to statistically significant dominance relationships between 12 of the 15 possible pairs of the six countries. Côte d'Ivoire dominates all other countries, followed by Togo, which dominates Benin, Burkina Faso and Niger. Benin and Mali also dominate Burkina Faso and Niger. Higher-order dominance tests cannot order Mali and Togo, but lead to Mali and Niger respectively dominating Benin in the second order and Burkina in the third order.

The results also translate into the finding of a considerable heterogeneity of the country-specific gaps between rural and urban poverty. The country rankings depend indeed considerably on whether we compare urban or rural poverty.

Burkina, which is found to be poorest when it comes to rural multidimensional poverty, exhibits lower urban poverty than Niger and Benin, and is then dominated only by Côte d'Ivoire. Benin is also inferred to be urban-wise poorer than all other countries. This suggests that it may be useful and informative to disaggregate multidimensional poverty comparisons within countries before proceeding to country-wise comparisons of welfare.

Table 1: Description of the Demographic and Health Surveys used

	Benin	Burkina	CI	Mali	Niger	Togo
Survey years	1996	1998-99	1998-99	1995-96	1998	1998
Number of households						
- Total	4499	4812	2122	8716	5928	7517
- Urban area (%)	32	26	67	32	28	32
- Rural area (%)	68	74	33	68	72	68
Number of women						
- Total	5491	6445	3040	9704	7577	8569
- Urban area (%)	33	26	68	36	31	36
- Rural area (%)	67	74	32	64	69	64
Number of men						
- Total	1535	2641	886	2474	3542	3819
- Urban area (%)	33	30	66	36	34	35
- Rural area (%)	67	70	34	64	66	65

Table 2: Descriptive statistics on the asset index (X)

Country	UBN procedure				MCA procedure			
	Mean	Std-dev.	Max	Min	Means	Std-dev.	Max	Min
Benin	0.009	0.573	1.930	-0.486	0.020	0.587	2.819	-0.456
Burkina	-0.150	0.472	2.045	$-0,486$	-0.133	0.496	2.819	-0.456
CI	0.663	0.754	2.044	-0.481	0.566	0.788	2.819	-0.449
Mali	0.066	0.498	1.989	-0.490	-0.010	0.496	2.647	-0.456
Niger	-0.179	0.482	2.009	-0.479	-0.159	0.485	2.819	-0.449
Togo	0.138	0.583	1.931	-0.491	0.104	0.580	2.819	-0.456
All	0.023	0.574	2.045	-0.491	0.000	0.572	2.819	-0.456

Table 3: Descriptive statistics on the nutritional indicator (Z_score)

Country	Mean	Std-dev.	Max	Min
Benin	-1.265	2.724	19.527	-28.471
Burkina	-1.749	2.512	17.108	-28.921
CI	-1.350	2.312	7.294	-27.808
Mali	-1.294	2.428	45.800	-23.978
Niger	-1.762	1.944	10.954	-13.117
Togo	-1.196	2.345	37.817	-29.016
All	-1.449	2.389	45.800	-29.016

Table 4: Sensitivity analysis on the asset index (X) — % of quartile not affording a good or a service

	UBN approach				MCA			
Quartiles of asset index	1^{st}	2^{nd}	3^{th}	4^{th}	1^{st}	2^{nd}	3^{th}	4^{th}
No electricity	100	100	100	57.3	100	100	99.7	57.0
No radio	82.7	31.6	35.2	18.9	72.0	56.2	20.7	16.7
No TV	100	100	100	59.0	100	100	99.8	58.5
No refrigerator	100	100	100	84.4	100	100	100	84.2
No bicycle	54.0	39.7	51.5	67.4	46.8	57.1	42.7	68.3
No motorbike	100	82.9	80.8	60.8	100	83.6	78.4	59.4
No car	100	100	99.9	87.4	100	100	99.9	87.2
Low quality floor	100	68.8	72.6	28.5	100	80.5	60.0	25.2
No toilet	100	95.0	36.0	14.8	100	78.4	45.5	18.2
No education	100	99.7	76.7	41.3	99.6	89.2	80.1	46.8
No water access	14.1	15.5	11.5	5.2	14.8	18.2	9.4	4.3

Table 5: The correlations between the asset indices produced by the UBN and MCA procedures

Country	Coefficients
Benin	0.97
Burkina	0.98
CI	0.94
Mali	0.96
Niger	0.97
Togo	0.96
All	0.96

Table 6: Tests for first-order stochastic dominance, country F versus country G, over a $[z^{-}, z^{+}]$ interval

Countries $F-G$	$[z^{-}, z^{+}]$ intervals Asset X	Z_score	p-value of rejecting non-dominance of F by G
Benin-CI	$[-0.46, 1.23]$	$[-2.96, 2.18]$	0.000^{***}
Burkina-CI	$[-0.48, 1.87]$	$[-5.97, 2.23]$	0.003^{***}
Mali-CI	$[-0.44, 1.52]$	$[-3.27, 2.30]$	0.000^{***}
Niger-CI	$[-0.45, 1.85]$	$[-3.46, 1.85]$	0.000^{***}
Togo-CI	$[-0.44, 1.00]$	$[-2.87, 2.18]$	0.000^{***}
Benin-Togo	$[-0.46, 0.24]$	$[-2.92, 2.36]$	0.090^{*}
Burkina-Togo	$[-0.47, 1.18]$	$[-5.64, 2.23]$	0.013^{**}
Mali-Togo	–	–	no dominance
Niger-Togo	$[-0.44, 1.63]$	$[-5.03, 1.97]$	0.033^{**}
Benin-Mali	–	–	no dominance
Burkina-Mali	$[-0.47, 1.08]$	$[-5.46, 2.44]$	0.000^{***}
Niger-Mali	$[-0.45, 1.62]$	$[-3.47, 2.28]$	0.000^{***}
Burkina-Benin	$[-0.47, 1.64]$	$[-5.85, 0.44]$	0.003^{***}
Niger-Benin	$[-0.44, 1.10]$	$[-3.38, 2.11]$	0.070^{*}
Burkina-Niger	–	–	no dominance

*: significant at 10%; **: significant at 5%; ***: significant at 1%.

Table 7: Tests of first-order stochastic dominance, rural area (F) versus urban area (G) within different countries, over a $[z^{-}, z^{+}]$ interval

Countries	$[z^{-}, z^{+}]$ intervals Asset X	Z_score	p-value of rejecting non-dominance of F by G
Benin	$[-0.48, 1.58]$	$[-3.00, 2.37]$	0.018^{**}
Burkina	$[-0.48, 1.68]$	$[-6.03, 2.43]$	0.028^{**}
CI	$[-0.40, 1.92]$	$[-2.86, 1.84]$	0.013^{**}
Mali	$[-0.46, 1.57]$	$[-5.02, 2.44]$	0.033^{**}
Niger	$[-0.45, 1.62]$	$[-4.88, 1.63]$	0.025^{**}
Togo	$[-0.47, 1.58]$	$[-5.08, 2.07]$	0.008^{***}
Total	$[-0.47, 1.73]$	$[-5.33, 2.27]$	0.000^{***}

*: significant at 10%; **: significant at 5%; ***: significant at 1%.

Table 8: Tests for higher-order stochastic dominance, country F versus country G, over a $[z^-, z^+]$ interval

Countries $F - G$	$[z^-, z^+]$ intervals		p-value of rejecting non-dominance of F by G
	Asset X	Z_score	
Benin-Mali (order 2)	$[-0.47, 0.10]$	$[-3.23, 2.71]$	0.020^{**}
Mali-Togo (orders 2 and 3)	–	–	no dominance
Burkina-Niger (3)	$[-0.45, 1.69]$	$[-5.44, -0.55]$	0.040^{**}

*: significant at 10%; **: significant at 5%; ***: significant at 1%.

Appendices

Appendix 1: Statistical inference for higher-order dominance

For a numerical solution to the problem of (18) with (19), (20) and $s_1, s_2 > 1$, consider the following Lagrangian ($\mathcal{L}$),

$$\mathcal{L} = \sum_i n_i^F \log P_i^F + \sum_j n_j^G \log P_j^G + \lambda_F(1 - \sum_i P_i^F) + \lambda_G(1 - \sum_j P_j^G)$$
$$- \mu(\sum_i P_i^F \Gamma_{F,i}^{s_1,s_2}() I_i(z_1, z_2) - \sum_j P_j^G \Gamma_{G,j}^{s_1,s_2}() I_j(z_1, z_2)), \tag{30}$$

with

$$\sum_i P_i^F \Gamma_{F,i}^{s_1,s_2}() I_i(z_1, z_2) = \sum_i P_i^F (z_1 - x_{i1}^F)^{s_1-1} (z_2 - x_{i2}^F)^{s_2-1} I\left(x_{i1}^F \leq z_1, x_{i2}^F \leq z_2\right)$$

and

$$\sum_j P_j^G \Gamma_{G,j}^{s_1,s_2}() I_j(z_1, z_2) = \sum_j P_j^G (z_1 - x_{j1}^G)^{s_1-1} (z_2 - x_{j2}^G)^{s_2-1} I\left(x_{j1}^G \leq z_1, x_{j2}^G \leq z_2\right)$$

and where λ_F, λ_G and $\mu \in R$ are Lagrange multipliers. The first-order conditions are given by:

$$\lambda_F + \lambda_G = N_F + N_G = N, \tag{31}$$

$$P_i^F = \frac{n_i^F}{\lambda + \mu \Gamma_{F,i}^{s_1,s_2}() I_i(z_1, z_2)} \quad \text{and} \quad P_j^G = \frac{n_j^G}{N - \lambda - \mu \Gamma_{G,j}^{s_1,s_2}() I_j(z_1, z_2)}, \tag{32}$$

with $\lambda = \lambda_F$. For given (z_1, z_2), it is then possible to solve the problem of maximizing (30) by searching for $\widehat{\lambda}$ and $\widehat{\mu}$ as follows:

$$(\widehat{\lambda}, \widehat{\mu}) = \arg\min_{\lambda,\mu \in R} - \sum_i n_i^F \log(\lambda + \mu \Gamma_{F,i}^{s_1,s_2}() I_i(z_1, z_2))$$
$$- \sum_j n_j^G \log(N - \lambda - \mu \Gamma_{G,j}^{s_1,s_2}() I_j(z_1, z_2)). \tag{33}$$

For all pairs of thresholds (z_1, z_2), the probabilities $\widehat{P}_i^F(z_1, z_2)$ and $\widehat{P}_j^G(z_1, z_2)$ are obtained by replacing λ and μ in (32) by their estimates $\widehat{\lambda}$ and $\widehat{\mu}$ at (z_1, z_2). The likelihood ratio is then given as

$$LR_{s_1,s_2} = 2 \left\{ \begin{array}{c} -N_F \log N_F - N_G \log N_G + \sum_i n_i^F \log n_i^F + \sum_j n_j^G \log n_j^G \\ -\sum_i n_i^F \log \widehat{P}_i^F(z_1, z_2) - \sum_j n_j^G \log \widehat{P}_j^G(z_1, z_2). \end{array} \right\} \tag{34}$$

We can then use the techniques of Davidson (2007) to show that this ratio is asymptotically equivalent to the square of a minimum t-statistic analogous to (29). The rest of the procedures is similar to those outlined on page 13 for first-order dominance.

Appendix 2: The Underlying Bivariate Normal (UBN) method

Consider the following:

$$p(f|x) = \frac{p(f)\Pr(x|f)}{\int_R p(f)\Pr(x|f)df}. \tag{35}$$

$p(f)$ can be derived from the assumption that the distribution of f is $N(0,1)$. Since the conditional distribution of x^* given f is $N(f\lambda, \psi)$, we have:

$$\Pr(x|f) = \Pr\left[\begin{array}{c|c} \gamma^1_{a_1-1} \le x^*_1 \le \gamma^1_{a_1} & \\ \cdot & \\ \cdot & f \\ \cdot & \\ \gamma^J_{a_J-1} \le x^*_J \le \gamma^J_{a_J} & \end{array}\right] =$$

$$\int_\Omega (2\pi)^{-J/2}|\psi|^{-1/2}\exp\left\{-\frac{(x^*-f\lambda)'\psi^{-1}(x^*-f\lambda)}{2}\right\}dx^*,$$

with

$$\Omega = \left[\begin{array}{c} \gamma^1_{a_1-1} \le x^*_1 \le \gamma^1_{a_1} \\ \cdot \\ \cdot \\ \cdot \\ \gamma^K_{a_K-1} \le x^*_K \le \gamma^K_{a_K} \end{array}\right], \tag{36}$$

$$p(f)\Pr(x|f) = (2\pi)^{-(J+1)/2}|\psi|^{-1/2}\int_\Omega \exp\left\{-\frac{g(x^*)}{2}\right\}dx^* \tag{37}$$

and where

$$g(x^*) = f^2 + (x^*-f\lambda)'\psi^{-1}(x^*-f\lambda) = B\left[f - \frac{1}{B}(\lambda'\psi^{-1}x^*)\right]^2 + x^{*\prime}Ax^*. \tag{38}$$

We also have that:

$$B = 1 + \lambda'\psi^{-1}\lambda \tag{39}$$

$$A = \psi^{-1} - \frac{1}{B}\psi^{-1}\lambda\lambda'\psi^{-1}. \tag{40}$$

When a quadratic loss function is used, the mean of the *a posteriori* distribution of the factor score is the bayesian estimator that minimizes the posterior expected loss. It is given by:

$$E\left(f|x\right) = \int_R f p(f|x) df \tag{41}$$

$$= \frac{\int_R f p(f) \Pr(x|f) df}{\int_R p(f) \Pr(x|f) df} \tag{42}$$

$$= \frac{\int_R f \int_\Omega \exp\left\{-\frac{g(x^*)}{2}\right\} dx^* df}{\int_R \int_\Omega \exp\left\{-\frac{g(x^*)}{2}\right\} dx^* df} \tag{43}$$

$$= \frac{\int_\Omega \int_R f \exp\left\{-\frac{g(x^*)}{2}\right\} df dx^*}{\int_\Omega \int_R \exp\left\{-\frac{g(x^*)}{2}\right\} df dx^*} \tag{44}$$

$$= \frac{\int_\Omega \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} \int_R f \exp\left\{-\frac{\left[f-\frac{1}{B}(\lambda'\psi^{-1}x^*)\right]^2}{2}\right\} df dx^*}{\int_\Omega \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} \int_R \exp\left\{-\frac{\left[f-\frac{1}{B}(\lambda'\psi^{-1}x^*)\right]^2}{2}\right\} df dx^*}. \tag{45}$$

Without loss of generality, moving the mean of f from 0 to $\frac{1}{B}\lambda'\psi^{-1}x^*$, it follows from the properties of the normal density function that:

$$\int_R \exp\left\{-\frac{\left[f - \frac{1}{B}(\lambda'\psi^{-1}x^*)\right]^2}{2}\right\} df = 1 \quad \text{and} \tag{46}$$

$$\int_R f \exp\left\{-\frac{\left[f - \frac{1}{B}(\lambda'\psi^{-1}x^*)\right]^2}{2}\right\} df = \frac{1}{B}\lambda'\psi^{-1}x^*. \tag{47}$$

This implies that

$$X_i = \widehat{f}_B = E\left(f|x\right) = \frac{\int_\Omega \frac{1}{B}\lambda' \psi^{-1} x^* \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} dx^*}{\int_\Omega \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} dx^*} = \frac{1}{B}\lambda'\psi^{-1}x_W^*, \quad (48)$$

with

$$x_W^* = \frac{\int_\Omega x^* \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} dx^*}{\int_\Omega \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} dx^*}. \quad (49)$$

Noting that $\exp\left\{x^{*\prime}Ax^*/2\right\}$ is proportional to the density function of the distribution $N(0, A^{-1})$, Shi and Lee (1997) suggest a simple Monte Carlo method by which L random vectors $u_1, ..., u_L$ are generated, where each vector is generated from a uniform distribution on Ω. Assuming that we have to estimate the integral $\int_\Omega Q(x^*)dx^*$, we have

$$V(\Omega)\left[Q(u_1) + ... + Q(u_L)\right]/L \to \int_\Omega Q(x^*)dx^* \;\text{ si }\; L \to \infty \quad (50)$$

where $V(\Omega)$ denotes the volume of Ω. By applying this result to x_W^*, we then find that

$$\frac{D(u_1) + ... + D(u_L)}{d(u_1) + ... + d(u_L)} \to x_W^*, \;\text{ si }\; L \to \infty \quad (51)$$

with

$$D(x^*) = x^* \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\} \;\text{ et }\; d(x^*) = \exp\left\{-\frac{x^{*\prime}Ax^*}{2}\right\}. \quad (52)$$

References

ANDERSON, G. (1996): "Nonparametric Tests of Stochastic Dominance in Income Distributions," *Econometrica*, 64, 1183–1193.

——— (2005): "Life Expectancy and Economic Welfare: The Example of africa in the 1990s," *Review of Income and Wealth*, 51, 455–468.

ATKINSON, A. B. (1987): "On the Measurement of Poverty," *Econometrica*, 55, 749–764.

——— (1998): *Poverty in Europe*, London: Blackwell.

——— (2003): "Multidimensional Deprivation: Contrasting Social Welfare and Counting Approaches," *Journal of Economic Inequality*, 1, 51–65.

BARRETT, G. F. AND S. G. DONALD (2003): "Consistent Tests for Stochastic Dominance," *Econometrica*, 71, 71–74.

BARTHOLOMEW, D. J. (1983): "Latent Variable Models for Ordered Categorical Data," *Journal of Econometrics*, 22, 229–243.

——— (1984): "The Foundations of Factor Analysis," *Biometrika*, 71, 221–232.

BELLMAN, R. (1961): *Adaptive Control Processes: A Guided Tour*, Princeton: Princeton University Press.

BOURGUIGNON, F. AND S. R. CHAKRAVARTY (2002): "Multidimensional Poverty Orderings," DELTA Working Paper 2002-22, DELTA.

CHEN, S. AND M. RAVALLION (2001): "How Did the World's Poorest Fare in the 1990s?" *Review of Income and Wealth*, 47, 283–300.

DAVIDSON, R. (2007): "Testing for Restricted Stochastic Dominance: Some Further Results," Technical report, McGill University.

DAVIDSON, R. AND J.-Y. DUCLOS (2000): "Statistical Inference for Stochastic Dominance and for the Measurement of Poverty and Inequality," *Econometrica*, 68, 1435–1464.

——— (2006): "Testing for Restricted Stochastic Dominance," IZA Discussion Paper No 2047, IZA.

DUCLOS, J.-Y., D. E. SAHN, AND S. D. YOUNGER (2006): "Robust Multimensional Poverty Comparison," *Economic Journal*, 113, 943–968.

FOSTER, J. E., J. GREER, AND E. THORBECKE (1984): "A Class of Decomposable Poverty Measures," *Econometrica*, 52, 761–766.

FOSTER, J. E. AND A. F. SHORROCKS (1988a): "Poverty Orderings," *Econometrica*, 56, 173–177.

——— (1988b): “Poverty Orderings and Welfare Dominance,” *Social Choice Welfare*, 5, 179–198.

GREENACRE, M. (1993): *Correspondence Analysis in Practice*, Academic Press.

GREENACRE, M. AND J. BLASIUS (2006): *Multiple Correspondence Analysis and Related Methods*, Statistics in the Social and Behavioral Sciences Series, CRC Press.

JÖRESKOG, K. G. AND I. MOUSTAKI (2001): “Factor Analysis of Ordinal Variables: A Comparison of Three Approaches,” *Multivariate Behavioral Research*, 36, 347–387.

KAUR, A., B. L. S. P. RAO, AND H. SINGH (1994): “Testing for Second-Order Stochastic Dominance of Two Distributions,” *Econometric Theory*, 10, 849–866.

KELLEY, A. (1991): “The Human Development Index: Handle with Care,” *Population and Development Review*, 17, 315–324.

MAASOUMI, E. (1999): “Multidimensional Approaches to Welfare Analysis,” in *Handbook of Income Inequality Measurement*, ed. by J. Silber, Boston: Kluwer Academic, 437–477.

MOUSTAKI, I. (2000): “A Latent Variable Model for Ordinal Variables,” *Applied Psychological Measurement*, 24, 211–223.

SAHN, D. E. AND D. C. STIFEL (2000): “Poverty Comparisons Over Time and Across Countries in Africa,” *World Development*, 28, 2123–2155.

——— (2002): “Robust Comparisons of Malnutrition in Developing Countries,” *American Journal of Agricultural Economics*, 84, 716–735.

——— (2003): “Exploring Alternative Measures of Welfare in the Absence of Expenditure Data,” *Review of Income and Wealth*, 49, 463–489.

SEN, A. (1985): *Commodities and Capabilities*, Amsterdam: North-Holland.

SHI, J.-Q. AND S.-Y. LEE (1997): “A Bayesian Estimation of Factor Score in Confirmatory Factor Model with Politomous, Censored or Truncated Data,” *Psychometrika*, 62, 29–50.

STREETEN, P., S. BURKI, M. UL HAQ, N. HICKS, AND F. STEWART (1981): *First Things First. Meeting Basic Human Needs in the Developing World*, New York and Oxford: World Bank and Oxford University Press.

UNITED NATIONS DEVELOPMENT PROGRAM (1990): *Human Development Report*, New York: Oxford University Press.