

WIDER Working Paper No. 2013/053

Advancing small area estimation

Channing Arndt¹, M. Azhar Hussain², Vincenzo Salvucci³,
Finn Tarp⁴ and Lars P. Østerdal⁵

May 2013

Abstract

The poverty mapping methodology for estimating welfare rankings from small areas has proven to be useful in guiding allocation of government funds, regional planning, and general policy formulation. Nevertheless, poverty mapping also suffers from a series of by now well recognized shortcomings. We apply an approach based on first order dominance (FOD) to small area estimation. Five advantages to the FOD approach are highlighted. First, it can serve as a complement to, substitute for, and/or extension of the poverty mapping methodology. Second, it directly uses census data with a minimum of assumptions imposed. Third, the methodology is straightforward to implement and the concepts are intuitive. Fourth, the FOD approach is multi-dimensional allowing for a broader conception of poverty. Finally, FOD indicators can be chosen that relate directly to public expenditure priorities...

Keywords: small area estimation, welfare, poverty mapping, multidimensional poverty measurement, first order dominance, Mozambique
JEL classification: C15, C81, I32

Copyright © UNU-WIDER 2013

¹Development Economics Research Group, University of Copenhagen, channing.arndt@econ.ku.dk; ²Roskilde University, Denmark, azharh@ruc.dk; ³Development Economics Research Group, University of Copenhagen, vincenzo.salvucci@gmail.com; ⁴UNU-WIDER, Helsinki, tarp@wider.unu.edu; ⁵University of Southern Denmark, lpro@sam.sdu.dk

This study has been prepared within the UNU-WIDER project 'Reconciling Africa's Growth, Poverty and Inequality Trends', directed by James Thurlow and Finn Tarp.

UNU-WIDER gratefully acknowledges the financial contributions to the research programme from the governments of Denmark, Finland, Sweden, and the United Kingdom.

ISSN 1798-7237

ISBN 978-92-9230-630-4



We apply the approach to census data from Mozambique for 1997 and 2007 and compare results with the poverty mapping methodology. We conclude that the FOD approach is well suited to small area estimation.

Acknowledgements

The authors wish to acknowledge the support of the National Directorate of Studies and Policy Analysis (DNEAP) within the Mozambican Ministry of Planning and Development and by UNU-WIDER. The support from DNEAP comes within the context of long-term institution-building efforts supported by the Danish International Development Agency and Swiss Development Cooperation.

The World Institute for Development Economics Research (WIDER) was established by the United Nations University (UNU) as its first research and training centre and started work in Helsinki, Finland in 1985. The Institute undertakes applied research and policy analysis on structural changes affecting the developing and transitional economies, provides a forum for the advocacy of policies leading to robust, equitable and environmentally sustainable growth, and promotes capacity strengthening and training in the field of economic and social policy making. Work is carried out by staff researchers and visiting scholars in Helsinki and through networks of collaborating scholars and institutions around the world.

www.wider.unu.edu

publications@wider.unu.edu

UNU World Institute for Development Economics Research (UNU-WIDER)
Katajanokanlaituri 6 B, 00160 Helsinki, Finland

Typescript prepared by Lorraine Telfer-Taivainen at UNU-WIDER

The views expressed in this publication are those of the author(s). Publication does not imply endorsement by the Institute or the United Nations University, nor by the programme/project sponsors, of any of the views expressed.

1 Introduction

The approach to ranking small areas in terms of their poverty levels for policy and other purposes has been the poverty mapping methodology, also called small area estimation (Elbers, Lanjouw and Lanjouw 2003; Tarozzi and Deaton 2009; Molina and Rao 2010). The methodology is now widely applied, particularly in developing countries, in order to produce information on welfare at small scales (such as district or even village). This information is useful to guide allocation of government funds, regional planning, and general policy formulation; and, as a consequence, the worldwide demand for poverty mapping is strong. While boasting considerable advantages, poverty mapping also suffers from a series of by now well recognized shortcomings. In this paper, we apply an approach based on first order dominance (FOD), recently developed by Arndt et al. (2012b), to small area estimation.¹

The FOD approach has a series of advantages in the context of small area estimation. First, it can serve as a complement to, substitute for, and/or extension of the poverty mapping methodology. As such, the approach either provides valuable information at small scales when the poverty mapping approach is not feasible or additional information when it is. Second, it directly uses census data with a minimum of assumptions imposed. Third, the methodology is straightforward to implement and the concepts are intuitive. This makes the approach particularly attractive in developing country settings. Fourth, the FOD approach is multi-dimensional allowing for a broader conception of poverty. Finally and importantly, FOD indicators can be chosen that relate directly to public expenditure priorities. If a welfare measure for small areas is going to help guide the allocation of public expenditures on items such as water, sanitation, education, and electrification across space, then direct indicators associated with these expenditure priorities would appear to be logical guides. The FOD permits the use of these indicators.

The remainder of this paper is structured as follows. Section 2 briefly presents the poverty mapping methodology including a discussion of the advantages/disadvantages associated with the approach. Section 3 presents the FOD methodology also including a discussion of the advantages/disadvantages of the approach. Sections 4 and 5 present an application to Mozambique. Section 4 covers the data and variables used while Section 5 details results including comparisons between poverty mapping and FOD. Section 6 concludes that the FOD approach provides a useful addition to the poverty measurement toolkit including application to small area estimation.

2 Poverty mapping methodology

The poverty mapping methodology is applied in cases where:

- i. There exists a survey with information on consumption (y) and household characteristics (X). Almost by definition, the survey is applied to a sample and does not cover all households in the targeted population. The ability to make viable inferences with respect to the welfare status of sub-populations in the country is determined by the sampling procedure.

¹ Here, we use the term ‘poverty mapping’ to refer to econometric methods like those of Elbers et al. (2003) and the term ‘small area estimation’ as a more general attempt to rank finely classified groups by welfare status.

- ii. There exists a census of the population that occurred in reasonable proximity in time to the household consumption survey. For practical reasons, censuses do not attempt to obtain consumption information from all households in the population. Instead, a census will often attempt to obtain information on household characteristics that are relatively easily observable for all (or a large subsample of) households in the targeted population. If the census and survey are designed with poverty mapping in mind, the collected household characteristics (X) will be comparable between the census and the survey.

When these two elements are present, the poverty mapping approach relies on a set of domain-specific survey-based regressions that model (per capita) log consumption, y , as a function of explanatory household and area level variables, X^{sur} , producing a vector of estimated parameters $\hat{\beta}^{sur}$, where superscript *sur* indicates that the variable/parameter is from the survey. These estimated parameters are combined with explanatory household and area level variables from the census (X^{cen}). By combining the estimated parameters from the survey and the household characteristics from the census, we are able to assign an expected household log consumption level $\hat{y} \left(= \hat{\beta}^{sur} X^{cen} \right)$ to each household in the census along with its estimated variance. Based on this information, we can estimate the probability of poverty for household i in a given small area:

$$(1) \quad \hat{h}_i = P(\hat{y}_i < \ln z) = \Phi\left(\frac{\ln z - \hat{y}_i}{\hat{\sigma}}\right)$$

where z is the poverty line and $\hat{\sigma}$ is the standard error of prediction from the regression in a given stratum and Φ is the cumulative distribution function of a standard normal. Small area poverty levels ($\hat{h}_j, j=1, \dots, N_j$) are then simply estimated as the average of household poverty probabilities weighted by the number of household members:

$$(2) \quad \hat{h}_j = \frac{1}{N_j} \sum_{i=1}^{N_j} \hat{h}_i w_i$$

where w_i is a weight proportional to the size of household i and N_j is the number of households in small area j .²

There is little doubt that small area estimates derived in the manner described have provided valuable information across a wide range of countries. The ubiquity of poverty mapping exercises attests to the demand for the information produced. Where possible to implement, these exercises should be conducted. Nevertheless, a series of shortcomings associated with the approach are by now well recognized. We list five shortcomings:

1. Most obviously, it is sometimes not possible to implement the poverty mapping approach even when census data are available. This occurs when a viable household survey, implemented in ‘reasonable’ proximity in time to the census, is not available. In principle and with patience, this issue can be addressed by assuring that the

² See Elbers et al. (2003) for details on the poverty mapping methodology including calculation of standard errors.

statistical authority conducts a household consumption survey within reasonable proximity in time to the next census.

2. Unfortunately, the existence of a timely household consumption survey only represents the beginning of a complex chain. The poverty mapping methodology is totally dependent on the measurement of real consumption in the survey as well as the estimation of poverty lines that reflect a reasonably constant living standard across space and time. This remains challenging and consumption poverty estimates from surveys are without doubt controversial. For example, in Tanzania, results from the 2007 Household Budget Survey (HBS) indicated a statistically insignificant decline in poverty compared with levels observed in 2000/01 with little change in inequality. This stagnation in poverty created considerable controversy because it was (and remains) difficult to reconcile with the economic growth rates reported by national accounts over the same period (Atkinson and Lugo 2010).

Further, in 2008/09, a new national panel survey (NPS) went into the field for the first time with funding from the Gates Foundation and analytical support from the World Bank. A second wave went into the field in 2010/11. These surveys contained consumption modules and poverty rates were estimated from this data. In 2008/09, the NPS measured the poverty headcount at 14.7 per cent while the 2007 HBS had measured it at 33.6 per cent. The wave 2 report states explicitly ‘the poverty analysis in this NPS report employs the same methodology as the HBS’ (National Bureau of Statistics 2012: 11). Thus, in principle, methodology does not account for the difference. Nor does anyone appear to believe that the poverty rate more than halved between 2007 and 2008/09. Instead, the wave 2 report points to ‘differences in the collection of consumption data in the NPS and the HBS’ (p. 11) as the source of difference.

It is certainly not uncommon for different data collection approaches to lead to considerable differences in poverty results (Deaton and Kozel 2005). At the same time, it is not appropriate simply to point to data collection differences and ignore the differences in results. After all, both surveys are, in principle, making a real attempt to measure *the same thing in the same country*. Nevertheless, the NPS poverty estimate is approximately 24 standard deviations from the HBS estimate using the published standard error on the 2007 poverty figure (National Bureau of Statistics 2009: 75).

Controversy over consumption poverty estimates is hardly confined to Tanzania. Grimm and Gunther (2007) report for Burkina Faso a decline in poverty of about 15 percentage points, from 62 to 47, over the period 1998-2003, while official estimates reported a one point rise over the same period. Deaton and Kozel (2005) review what they term ‘the great Indian poverty debate.’ Similar to Tanzania, they report that a shift in reporting period in the consumption questionnaire generated gains in measured consumption sufficient to halve the poverty rate. Recently, Alfani et al. (2012) argue that poverty in Mozambique declined by four percentage points between 2002/03 and 2008/09 as opposed to remaining essentially constant as indicated in official figures.

Of course, controversies extend beyond the national numbers and into the regional distribution and poverty profile as well. The point of all this is that appropriate measurement of consumption poverty is difficult and that at least some controversy

over consumption poverty measurements exists in nearly every country. Given the complete dependence of poverty mapping on the quality of estimates from the household survey, these controversies and differences in results translate over to the poverty mapping results.

3. Even if one is fully comfortable with the household consumption survey and associated poverty estimates, it is almost surely the case that the poverty estimates in the household survey reflect a particular conjuncture of events. Poor households, particularly those in poor societies, frequently lack the means to substantially smooth consumption in the face of shocks. For example, Grimm and Gunther (2007) attribute part of the 15 point decline in poverty they observed in a five year period in Burkina Faso to drought conditions experienced in the initial period. Of course, events vary across space as well. For Mozambique, Alfani et al. (2012) report a 25 percentage point decline in poverty in the rural zones of Niassa and Cabo Delgado provinces from 2002/03 to 2008/09 and a 12 percentage point increase in poverty in the rural areas of Sofala and Zambezia provinces over the same period. The official changes in poverty for these regions are very similar for the same time periods (DNEAP 2010).

These changes in poverty represent a difficult to decompose mix of sample error, non-sample error, and actual changes in living standards, which, in turn, are almost surely driven by a combination of underlying development progress (or lack thereof) and ephemeral shocks to welfare that may be either positive or negative. Correlation of measured changes in poverty with other observables can help to raise confidence that one is observing actual changes rather than statistical noise. For example, Arndt et al. (2012a) observe strong correlations between relative price changes for food, as recorded by the market information system, and observed changes in poverty in Mozambique between 2002/03 and 2008/09 pointing to food supply conditions as a strong determinant of poverty. In the regressions required to predict consumption poverty for poverty mapping purposes, it would be no small feat to consign the implications of these shocks (along with the sample and non-sample error) strictly to the error term arriving at a predicted consumption level under ‘normal’ conditions. In summary, the regressions necessary for poverty mapping are far from simple to implement and are quite likely to reflect patterns of poverty driven by the conjuncture of shocks that characterized the period of the survey.

4. An important assumption behind the poverty mapping methodology is that the domain is acceptably homogenous such that the consumption regression is applicable to the small areas within the domain. Also, area level regressors should capture sub-domain spatial correlation. Tarozzi and Deaton (2009) point out that even small violations of the ‘area homogeneity’ assumption may result in misleading inference. Tarozzi and Deaton also express concerns that sub-domain spatial integration might not be completely taken into account which leads to an underestimation of standard errors of small area estimates.³

Point estimates can also be influenced. In consumption regressions, unobserved heterogeneity across survey strata (e.g., across provinces) is often absorbed through dummy variables to capture domain specificity. As these average characteristics of the

³ In an evaluation of the poverty mapping methodology conducted on Brazilian data, both of these concerns seemed to be of minor importance (Elbers, Lanjouw and Leite 2008).

domain are then applied to all small areas within the domain, this can cause similar small areas separated by a border between two strata to be accorded different welfare rankings. These observations further underline the potential complexity of the consumption regressions.

5. Finally, since the small area estimates are fundamentally based on consumption, they still only provide a one-dimensional measure of poverty even though poverty has long been recognized as a multidimensional phenomenon.

These shortcomings do not combine to render the poverty mapping approach uninformative. Rather, they point to the utility of an expanded toolkit that can complement the poverty mapping methodology and provide rigorous welfare rankings across space and through time based on census data.

3 FOD Methodology

Arndt et al. (2012b) introduce the basic FOD criterion in the context of population welfare comparisons and discuss the advantages of the FOD approach compared with the (large) literature on multi-dimensional poverty/welfare measurement. The FOD criterion corresponds to what in probability theory is simply referred to as the ‘usual (stochastic) order’, cf., e.g., Lehmann (1955).⁴ In this section, we are principally concerned with the application of the FOD approach to small area estimation. After briefly touching on FOD in the context of other multidimensional approaches, we provide an intuitive and a mathematical discussion of the FOD approach. Next, we consider the application of the FOD approach in the small area estimation context.

Regarding multidimensional welfare comparisons, we refer to the literature on ‘robust’ methods for comparing population welfare, poverty and inequality.⁵ These methods rely on stochastic dominance concepts for comparisons that are valid for broad classes of underlying social welfare functions. As detailed in Arndt et al. (2012b), a principal difference and advantage of the FOD approach compared to other multidimensional approaches is that the FOD does not depend on a weighting scheme or on *ad hoc* simplifying assumptions about the social welfare function. Instead, for the case of binary welfare indicators where individuals or households are either deprived or non-deprived in a specific welfare dimension, the FOD criterion simply asserts that it is better to be non-deprived than deprived in any given dimension.

Intuition into the FOD approach is best gained by example. Suppose that we have data for five binary welfare indicators on populations A and B, and we wish to determine whether population A is unambiguously better off than population B based on these indicators. The respective populations can be divided into $2^5=32$ states corresponding to whether they are deprived or not deprived in the various dimensions. Obviously, if being not deprived is better than being deprived, then those who are not deprived in any dimension are best off and those deprived in all dimensions are worst off. Intermediate rankings are somewhat more complex.

⁴ For a general treatment of stochastic dominance theory, we refer to Müller and Stoyan (2002), or Shaked and Shanthikumar (2007).

⁵ E.g., Atkinson and Bourguignon (1982); Duclos, Sahn, and Younger (2007); Bourguignon and Chakravarty (2003); Batana and Duclos (2010).

If we define 0 as deprived and 1 as not deprived, then the state (0,1,1,0,0) is unambiguously better than (0,0,1,0,0) because the former state is always at least equivalent and is better than the latter in one instance. However, the states (1,0,1,0,0) and (1,1,0,0,0) are indeterminate (without further assumptions) because each state is better than the other in one dimension. The FOD criterion is strict. The state (1,1,0,1,1) is not unambiguously better than the state (0,0,1,0,0) because no judgment is made as to the relative importance of dimension three versus all other dimensions.

Formally, population A first order dominates population B if one can generate the shares of the population in each state in population B by shifting probability mass within population A to states that are unambiguously worse.⁶ Helpfully, this condition can be defined as a transportation problem with limitations on allowed transfer paths (e.g., Preston 1974).

In the applications considered below, we retain five dimensions of welfare or $2^5=32$ states. Building on Arndt et al. (2012b, Appendix A), we present the transportation problem for the five dimensional case. Define binary indices i, j, k, l, m , which each can take the value 0 or 1. The value 1 refers to not deprived and the value 0 to deprived for the five dimensions. Define binary indices i', j', k', l', m' , which are aliases of i, j, k, l and m respectively. For the two populations A and B, let a_{ijklm} and b_{ijklm} be the shares of the respective populations corresponding to the state of deprived and not deprived for the five indicators. Define the variable $x_{ijklm, i'j'k'l'm'}$ which represents transfer of probability mass from state $(ijklm)$ to state $(i'j'k'l'm')$. Define Z as the set of source-destination pairs $(ijklm, i'j'k'l'm')$ that move probability from preferred to less preferred states. If state $(ijklm)$ is the source of the transfer and state $(i'j'k'l'm')$ is the destination, a legal transfer is where $i' \leq i, j' \leq j, k' \leq k, l' \leq l$, and $m' \leq m$. Under these conditions, population A FOD population B if and only if the following linear program is feasible.

Min $y = 1$

subject to:

$$\begin{aligned}
 a_{ijklm} + \sum_{(i'j'k'l'm',ijklm) \in Z} x_{i'j'k'l'm',ijklm} \\
 - \sum_{(ijklm,i'j'k'l'm') \in Z} x_{ijklm,i'j'k'l'm'} &= b_{ijklm} \quad \forall i, j, k, l, m \\
 x_{ijklm, i'j'k'l'm'} \geq 0, \quad x_{ijklm,ijklm} &= 0.
 \end{aligned}$$

With the ability to compare any two populations, large numbers of population comparisons are possible. Suppose a census contains five binary welfare indicators of interest and an adequate number of observations for 100 distinct regions. The door is then open to running $100^2-100=9,900$ comparisons. Defensible welfare rankings of regions can be generated by, for example, counting the number of times a given region dominates all other regions and subtracting the number of times the same region is dominated by other regions generating a

⁶ This is equivalent to the condition that population A has higher welfare than population B for *any* increasing social welfare function; see Strassen (1965); Levhari et al. (1975), Grant et al. (1992); see also Østerdal (2010).

score in the interval $[-99,99]$. Regions can then be naturally ranked with higher scores superior to lower scores.

Note that if region A dominates region B and region B dominates region C, then region A must dominate region C.⁷ Therefore, if A dominates B, then the score of A must be strictly greater than the score of B using the above scoring approach (unless A and B are equivalent). The large number of comparisons inherent in small area estimation helps to mitigate the two principal disadvantages of the FOD approach. These two disadvantages, discussed in Arndt et al. (2012b), are:

- i. FOD comparisons are frequently indeterminate. If neither A dominates B and nor B dominates A (e.g., both linear programs are infeasible), then the welfare rankings of A and B are indeterminate without further information.
- ii. Suppose that A dominates B. The degree of domination is unknown without further information. The conclusion of dominance may rest on fine differences in the distributions across states or A could comfortably dominate B.

The multiple comparisons inherent in small area estimation generate additional information that offsets these disadvantages. Suppose that neither A nor B dominates the other but on net A dominates 20 other regions while B dominates negative one (e.g., the total number of regions that dominate B is one larger than the number of regions that B dominates). It is then sensible to rank A above B. The same logic applies to degree of dominance. The suggested scoring system naturally incorporates this additional information.

Arndt et al. (2012b) apply the logic of generating information through large numbers of comparisons to a setting where survey (rather than census) information is available. Surveys are normally representative of only a limited number of regions within a country. For example, surveys in Mozambique typically only define 11 statistically valid sub-regions (provinces) generating only 110 comparison pairs. Hence, the quantity of additional information generated by complete enumeration of all possible (statistically valid) inter-regional comparison combinations is relatively small. To mitigate this, Arndt et al. (2012b) apply a bootstrap technique. In addition to comparing the point estimates of distributions A and B, they generate K distribution pairs using the bootstrap technique and apply the FOD criteria to all pairs.

The application of the bootstrap is natural in the context of a survey but less so in the context of a census. Nevertheless, there are a number of reasons to apply the bootstrap to small area estimation using FOD. First, many countries have censuses from multiple years. It is then pertinent to use the FOD criterion to determine whether welfare has been improving through time at various geographical scales. The bootstrap then effectively mitigates the two disadvantages associated with the FOD approach while focusing on exactly the comparison of interest (region D at time t versus region D at time $t+1$). Second, with respect to spatial comparisons, while the marginal gain in information generated by bootstrapping regional comparisons is perhaps not as important as in the temporal case due to the already large

⁷ This is so because it is possible, by dominance of A over B, to generate the distribution of region B by shifting mass towards unambiguously worse outcomes starting with the original distribution of region A. By definition of B dominating C, the process of shifting mass within A can then continue until distribution A generates distribution C purely by shifting mass to worse outcomes proving that A dominates C.

numbers of comparisons normally inherent in small area estimation, the potential information gain can be obtained at relatively low cost-- some additional computer time and data management is all that is required. Finally, while bootstrapping census data is not normally done, the principles and concepts remain straightforward. Indeed, bootstrapping the census is simply random sampling from a population with replacement.

4 Data and empirical choices

The datasets used in the study are the 1997 and 2007 censuses, and the consumption surveys IAF 1996/97 (*Inquérito aos Agregados Familiares 1996/97*) and IOF 2008/09 (*Inquérito aos Agregados Familiares sobre Orçamento Familiar 2008/09*) from Mozambique. The surveys and censuses are conducted by the National Statistical Institute (INE). The surveys include information on general characteristics of the individuals and of the households, and there is information on daily, monthly and own consumption, and also information on possession of durable goods, transfers and gifts. The surveys have been used to estimate the first and also the latest set of poverty rates at the national and regional levels. Details on the poverty calculations and supplementary information for the IAF 1996/97 can be found in DNPO (1998) while for the IOF 2008/09 it can be found in INE (2010) and DNEAP (2010).

4.1 Variables used in the poverty mapping

The information included in the poverty mapping analysis is limited by what is available in both the census and in the survey. The common information available covers demographic characteristics, education, assets, own production of food items, and labour market variables. The same set of candidate variables are applied for 1996-97/1997 and 2008-09/2007. One area level variable is also included. It is a composite index made up of the average fraction of the population with certain characteristics assumed to influence consumption levels. This includes (fraction of) male-headed households, number of people aged 15-64 years, one minus the dependency ratio, different educational levels, own production of food items, economic activity, and non-disability. Consumption data are corrected for underreporting of calorie intake in specific regions. Details on the correction procedure can be found in DNEAP (2010).

4.2 FOD indicators

Five welfare indicators are considered, inspired by the notion of severe deprivation based on the Bristol indicators (Gordon et al. 2003). For the first two indicators, safe water and sanitation, the 2007 census questionnaire is more elaborate than the 1997 version allowing a more refined definition of deprived versus not deprived. In order to profit from the enhanced specificity in 2007, the definitions of deprived versus not deprived differ slightly between 1997 and 2007 for the spatial analyses (within year comparisons). Of course, for the temporal analysis, the coarser definitions from 1997 must also be applied in 2007. The indicators are:

Safe water

For 1997, there is access to safe water (not deprived) when the water source is piped water inside or outside the house or the water source is standpipes. For 2007, the water source should be piped water inside or outside the house/yard, spring water, hand pumped well water, or mineral/bottled water.

Sanitation

For 1997, we define the household as having access to sanitation (not deprived) when there is a flush toilet or a latrine. For 2007, we define it as having access to flush toilet, toilet with septic tank, or an improved latrine.

Education

This indicator takes the value 1 (not deprived) for households where at least one household member has some education.

Electricity

This indicator takes the value 1 (not deprived) for households with electricity for lighting.

Radio

This indicator takes the value 1 (not deprived) for households with a functioning radio.

Descriptive statistics for each welfare indicator are presented in Table 1.

Table 1: Descriptive statistics for district welfare indicators, 1997 and 2007 (%)

Indicator	1997				2007			
	Mean	St. Dev.	Min	Max	Mean	St. Dev.	Min	Max
Water	13.1	22.0	0	99	18.9	18.6	0	89
Sanitation	29.5	28.5	0	98	48.1	21.5	6	92
Education	65.8	17.0	26	99	84.7	9.3	56	100
Electricity	4.7	10.6	0	80	8.7	16.3	0	96
Radio	31.8	15.2	9	84	49.0	19.5	0	85
Non-poverty	33.1	15.0	4	84	52.5	16.9	15	98

Note: $N = 146$ (districts).

Source: authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/97 and IOF 2008/09 from Mozambique.

4.3 FOD metric: spatial case

One hundred and forty-six districts are identified in both 1997 and 2007, making for 21,170 total comparisons for each year without the bootstrap. To summarize these FOD spatial analyses for 1997 and 2007, we define a measure of dominance labelled 'net spatial domination', which is essentially a scaled version of the index discussed in section II. This is defined, for each district i , as the percentage of districts that are dominated by i minus the percentage of districts that dominate i . Hence, higher values in the net spatial domination index indicate that the district is relatively better off. This measure lies in the interval -1 to 1, making it somewhat less convenient to compare with the headcount ratio, which lies in the [0,1] interval. To aid visual presentation, we transform the net spatial domination index to lie in the [0,1] interval with higher values corresponding to greater deprivation. The

transformation implemented is: $spatial\ FOD\ index = \frac{1 - net\ spatial\ domination}{2}$.

The bootstrap approach discussed in section II was also employed for the analysis. The size of each bootstrap sample was chosen to be equal to the number of households in the least populous district, 1,828 households. One hundred bootstrap repetitions were employed generating more than two million potential spatial comparisons for each census year. Due to the reasonably large number of comparisons conducted without the bootstrap, net domination measures generated with and without the bootstrap give very similar results. We opt to present spatial results without the bootstrap.

4.4 FOD metric: temporal case

In the temporal case, we analyse for each district whether the 2007 welfare distribution dominates the welfare distribution of the same district in 1997, or whether 2007 is dominated by 1997. For each district, we define three possible results:

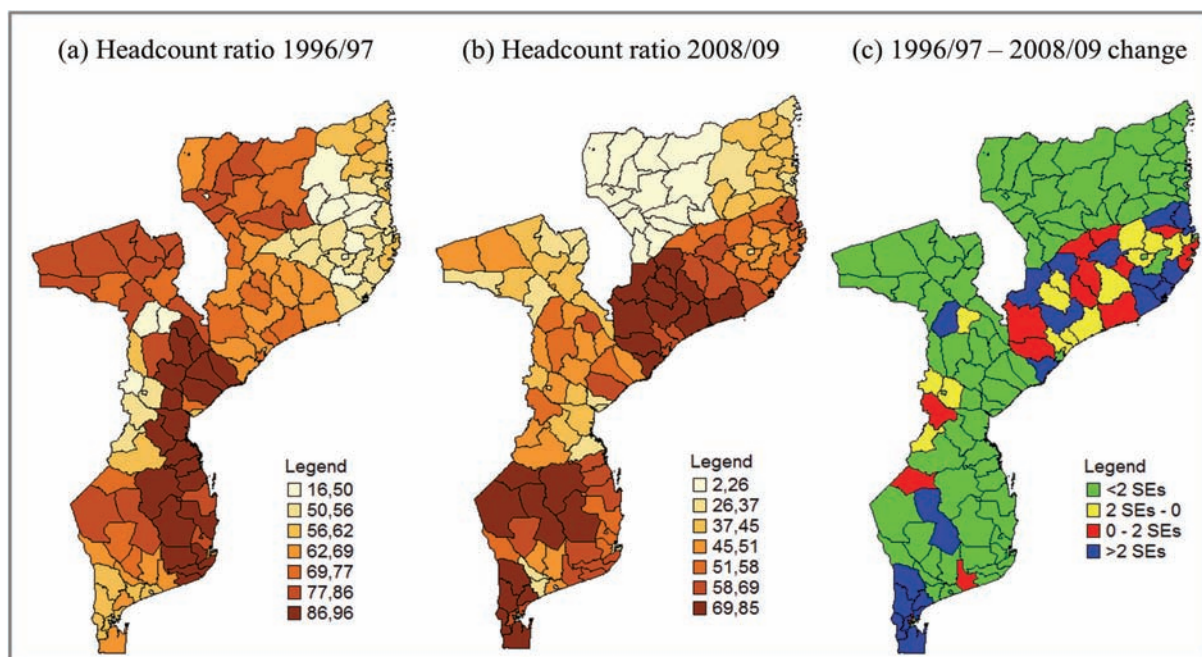
- 0: neither 2007 FOD 1997 nor 1997 FOD 2007
- 1: 2007 FOD 1997
- 1: 1997 FOD 2007.

We apply the bootstrap to the temporal case in order to generate probabilistic measures of dominance. The bootstrap sample is 1,828 households drawn 100 times. In the event, 1997 never dominates 2007 for any bootstrap draw across all regions. Consequently, simple averaging across the outcomes, either a zero or one, generates a probability of temporal domination. This probability is used as a measure of domination and we call it the ‘temporal FOD index’.

5 Results

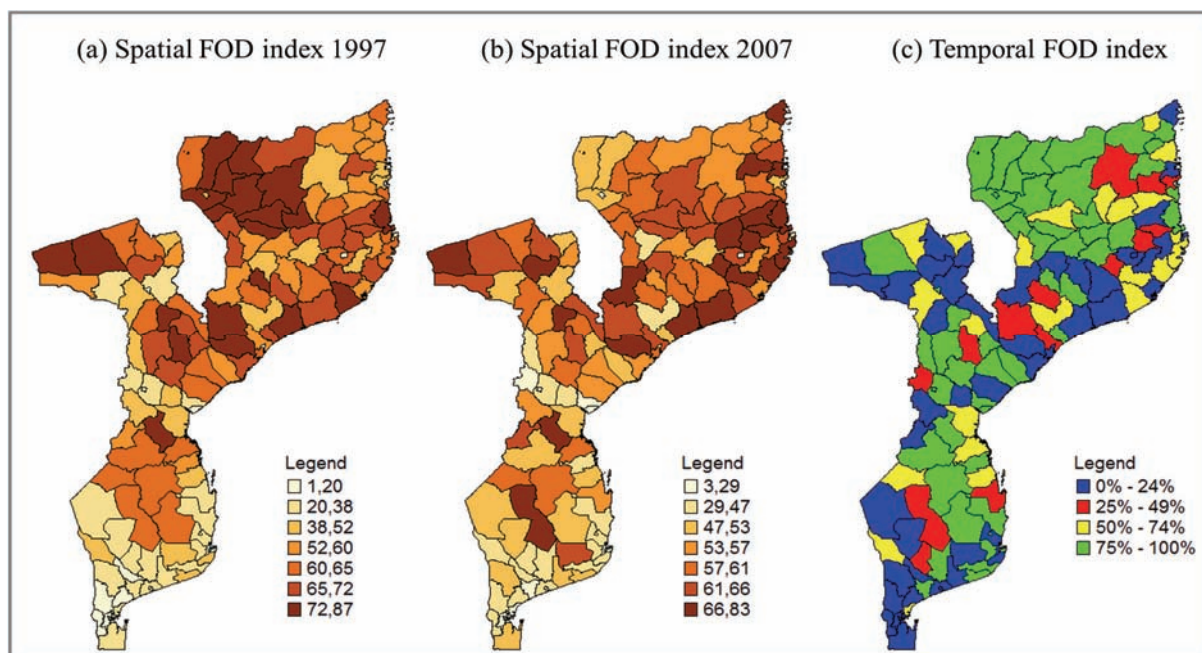
Figure 1 illustrates the poverty mapping results for 1996/97 (panel *a*), for 2008/09 (panel *b*), and for the change in the headcount ratio between 1996/97 and 2008/09 (panel *c*). Similarly, in Figure 2, the spatial FOD index for the two years 1997 and 2007 (panels *a* and *b*) and the temporal FOD index (panel *c*) are presented. Note that the poverty map refers to the share of the population living below some absolute welfare cut-off. Hence, the levels in panels (*a*) and (*b*) are comparable as these levels are, in principle, both relative to a fixed reference point. In contrast, for the FOD, the index levels registered in panels (*a*) and (*b*) are not comparable because they are respectively relative to the situations prevailing in 1997 and 2007. The temporal FOD index provides a measure of change through time. Descriptive statistics for the indices used are displayed in Table 2.

Figure 1: Poverty mapping, 1996/97 and 2008/09 (%)



Source: authors' analysis based on the consumption surveys IAF 1996/97 and IOF 2008/09 from Mozambique.

Figure 2: FOD mapping, 1997 and 2007



Note: In panels a and b, lower numbers represent superior district rankings, while in panel c higher numbers represent higher probability of progress.

Source: Authors' analysis based on the 1997 and 2007 censuses from Mozambique.

Table 2: Descriptive statistics for the indices used

Index and year	Mean	St. Dev.	Median	Min	Max
Headcount ratio 1996/97	66.9	15.0	65.1	16.4	96.0
Headcount ratio 2008/09	47.5	16.9	49.0	2.3	84.8
Headcount change, 1996/97- 2008/09	-19.3	20.9	-17.5	-61.1	20.7
Spatial FOD index 1997	50.0	21.8	57.6	1.0	86.9
Spatial FOD index 2007	50.0	18.3	55.0	2.8	83.4
Temporal FOD index	48.6	40.4	56.0	0.0	100.0

Note: $N = 146$ (districts).

Source: authors' analysis based on the 1997 and 2007 censuses, and the consumption surveys IAF 1996/97 and IOF 2008/09 from Mozambique.

5.1 Poverty mapping results

The 146 districts in Figure 1 are coloured depending on seven ordered levels of the headcount ratio. The levels are chosen such that (roughly) an equal number of districts are in each level.⁸ For 1997, we see that the highest poverty levels are found in the coastal zones of the centre-south, while the least poor districts are located in the south, close to the capital. Disaggregating the analyses to the district level shows that relatively richer provinces also have pockets of districts with high poverty rates.

In 2008/09, the poverty map changes. Of the districts placed in the poorest of the seven categories in 1997, only two districts remain among the most poor. Districts ranked as among the poorest now appear in the south, excluding the capital Maputo, and in the central province of Zambezia. We see that, in 2008/09, the northern and western parts of Mozambique, many districts in central Mozambique and a few north-eastern and southern districts are in the least poor groups.

An overview of districts' poverty trends from 1996/97 to 2008/09 is shown in panel *c* of Figure 1. Here the districts in which consumption poverty decreased more are marked in green, those in which it was reduced but the decrease was less than two standard errors are in yellow. The districts in which consumption poverty increased but less than two standard errors are marked in red, while districts in which it increased by more than two standard errors are in blue. Looking at the district level poverty change, 77 per cent of the districts experienced consumption poverty reduction over the decade from 1996/97 to 2008/09. On average, the reduction among those districts with falling poverty was 27 percentage points. Most of the districts with the largest poverty reductions are in the north-western province of Niassa. On the other hand, 23 per cent of the districts saw an increase in poverty, with an average increase of 5 percentage points. Districts with the largest poverty increases (above 10 percentage points) were all, but one, located in the southern Maputo province.

In sum, the change in consumption poverty over the decade has not been uniform. Rather, there is a tendency that districts with initially high poverty rates experienced the largest poverty reductions (panel *c* in Figure 1). This means that, on average, districts with high poverty rates in 1996/97 saw the greatest reductions in poverty rates. In this sense, the

⁸ In Figure 1, we are more interested in the relative rankings of the districts in each year, so we do not consider the same intervals of 1996/97 for the 2008/09 poverty map.

reduction in consumption poverty over 1996/97 to 2008/09 is characterized as having a pro poor bias.

5.2 FOD Results

The 146 districts in Figure 2 showing the FOD results are also coloured depending on seven ordered levels, depending on the previously introduced spatial FOD index. Also in this case, the levels are chosen such that (roughly) an equal number of districts are in each level. In panel *a* of Figure 2 (corresponding to 1997), we see that the districts ranked as most deprived are those located in the northern and central areas. None of the southern districts appear in the worst-ranked group. As it emerged in the headcount ratio analysis (poverty mapping), the FOD results also show that disaggregating the analysis to the district level provides additional information on intra-province welfare differences. In 2007 (Figure 2, panel *b*), we see that most of the FOD worst-ranked districts are again located in the central and northern provinces, while most of the southern districts are confirmed as the FOD best-ranked ones.

Looking at the FOD temporal index, as indicated, no district in 2007 is dominated by itself in 1997. There are 76 districts (out of 146) for which the probability of experiencing welfare improvement is higher than 50 per cent (Figure 2, panel *c*). Given the strictness of the FOD criterion, this is a salient result confirming broad based advance in living conditions between 1997 and 2007.

5.3 Poverty mapping versus FOD mapping

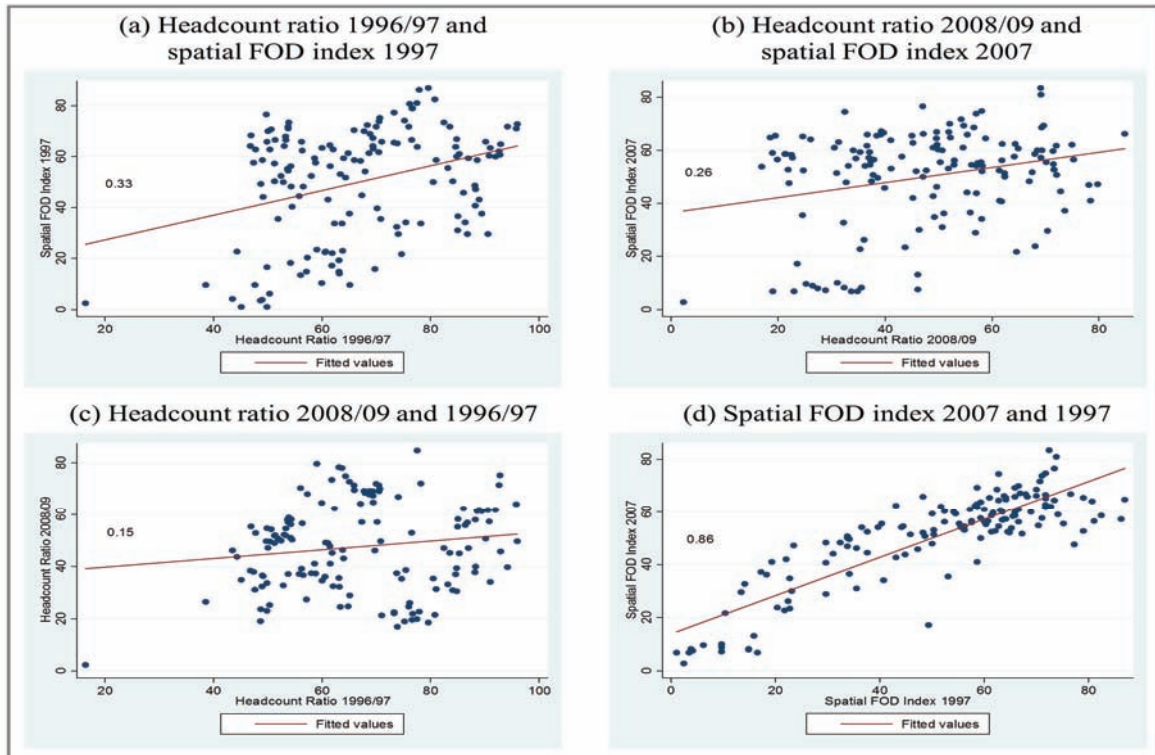
Comparing the poverty mapping and the FOD small area results for 1996-97/1997, we find that the rankings of poorer and richer districts are sensibly different: 47 districts change their ranking by more than 50 positions; and in particular most centre-northern districts are much better ranked in the poverty mapping than in the FOD mapping, while the opposite holds for the majority of districts in the centre-south. Similar results prevail when comparing the poverty mapping for 2008/09 and the FOD mapping in 2007. In 2007, the two most northern provinces score better in the poverty mapping than in the FOD analysis, while the three southern provinces (excluding Maputo City) are much better ranked in the multidimensional FOD than in the poverty mapping.

We undertake a correlation analysis based on the rankings obtained so as to provide a finer overview of the differences between poverty mapping and FOD. In Figure 3, scatter plots of the 146 Mozambican districts for 1996-97/1997 (panel *a*) and 2008-09/2007 (panel *b*) are displayed. On the horizontal axis we show the district headcount ratio (1996/97 and 2008/09) and on the vertical axis the corresponding spatial FOD index (1997 and 2007). The correlation coefficient between the two indices is 0.33 for 1996-97/1997 and 0.26 for 2008-09/2007. Concerning the changes over the analysed decade (not shown in the graph), a correlation of 0.33 is observed between the two welfare indices, suggesting that the two methodologies both capture a positive trend.

The lack of strong correlations suggests, at least in part, that different dimensions of welfare are being measured with the two approaches. In particular, the poverty mapping measure, effectively based on consumption, is strongly influenced by food availability as proxied by relative prices (Arndt et al. 2012a). Variations in food prices and availability can generate strong changes in consumption poverty measures. These strong variations are reflected in the

official poverty measures (DNEAP 2010) as well as in Alfani et al. (2012). Sample and non-sample error also undoubtedly contribute to the re-rankings. In contrast, the five indicators underlying the FOD indices tend to be a lot less volatile than consumption. Only the presence of a functioning radio would plausibly vary substantially with, for example, the quality of the agricultural season. In addition, a census is not subject to sample error. Non-sample error is present in every census/survey; however, the five indicators underlying the FOD are relatively simple to observe, especially in comparison with per capita household consumption, and thus less subject to non-sample error.

Figure 3: Correlation between headcount and FOD measures



Note: The correlation coefficient is presented on the left-hand side of each figure.

Source: authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/97 and IOF 2008/09 from Mozambique.

Volatility in the consumption measure and relative stability of the FOD indices are reflected in the correlations between the headcount from 1996/97 and 2008/09 and between the spatial FOD indices for 1997 and 2007. Indeed, the correlation coefficient between the headcount ratio in 1996/97 and in 2008/09 is low at 0.15 (Figure 3, panel c). Conversely, the correlation coefficient for the spatial FOD index in 1997 and in 2007 is relatively high at 0.86 (Figure 3, panel d). In sum, district welfare rankings are substantially more stable over time when based on the FOD welfare approach.

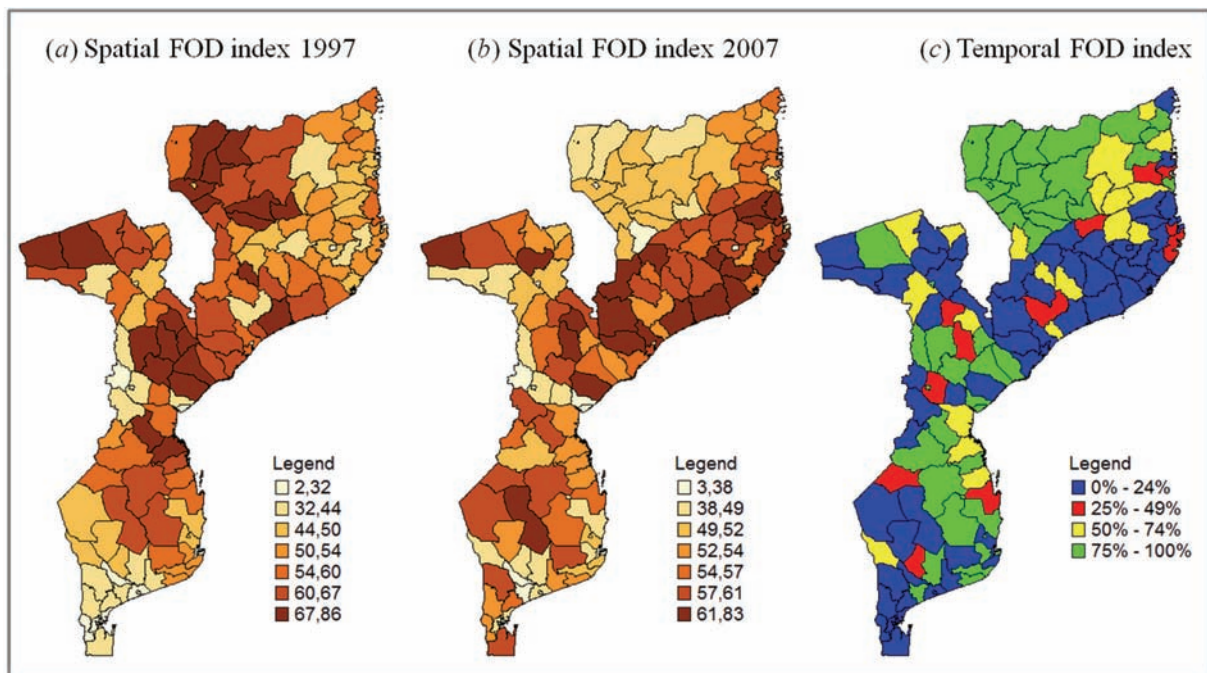
Finally, the effects of accounting for domain specificity in the consumption regressions are fairly clear from Figure 1. Even without prior knowledge of provincial administrative boundaries in Mozambique, a detailed look at the three panels of Figure 1 would provide solid hints as to the locations of at least some provincial boundaries. This is mainly an artefact of the inherent difficulties in using results from a sample, which in this case is designed to provide averages by province, to estimate welfare levels in all districts for all

provinces. For the FOD, provincial boundaries are irrelevant to the calculations. Correspondingly, the implications of provincial boundaries are a lot less marked in Figure 2.

5.4 Poverty as indicator instead of radio in FOD mapping

Here, we employ the FOD as an extension of the poverty mapping methodology by substituting the poor/non-poor indicator from the poverty mapping analysis for the radio indicator among the five FOD variables. Not surprisingly, when consumption poverty is included, the FOD mapping produced appears to be more similar to the results derived from the poverty mapping approach. Nonetheless, the rankings obtained from the FOD with consumption poverty rather than radio as a welfare indicator do not differ very much from those generated in the base case: the correlation being 0.84 for 1997 and 0.80 for 2007. The correlation for the temporal FOD index in the two cases is slightly higher (0.88). Figure 4 illustrates the spatial FOD index for 1997 and 2007 (panels *a* and *b*) and the temporal FOD index (panel *c*) when consumption poverty is taken as a welfare indicator.

Figure 4: FOD mapping with non-poverty from poverty mapping as a welfare indicator, 1997 and 2007



Source: authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/97 and IOF 2008/09 from Mozambique.

6 Conclusions

We conclude that the FOD approach to small area estimation has considerable advantages and represents a useful addition to the welfare analysis toolkit. The approach is flexible, robust, intuitive, straightforward to apply and multi-dimensional. Due to the large number of comparisons inherent in small area estimation, it appears to be particularly well suited to this class of problem. The direct use of multiple indicators from census data constitutes a further significant advantage.

For the case of Mozambique, the approach confirms broad based progress across a number of welfare indicators for the majority of districts. Importantly, there is no evidence of regress in any district and positive probability of progress in nearly all districts. In addition to measuring progress through time, district welfare rankings are obtained for 1997 and 2007. Because the rankings are based on relatively stable and easy to observe indicators, the rankings themselves are relatively stable through time, especially when compared with traditional poverty mapping. Also, four of the five indicators employed to develop the rankings relate directly to priority public expenditures in water, sanitation, education, and infrastructure (electricity). Consequently, these rankings would appear to form a reasonable basis for the allocation of public funds across districts.

With respect to future research, the FOD approach opens the door to robust welfare comparisons across a vast array of populations. These populations could be countries, ethnic groups, age groups and other criteria and combinations. As always, valid and comparable indicators are required. Exploitation of large and explicitly cross country data sets, such as the Afro-Barometer, may be a useful way forward.

References

- Alfani F., Azzarri C., d'Errico M. and Molini V. (2012). 'Poverty in Mozambique. New evidence from recent household surveys'. Policy Research Working Paper 6217. Washington DC: World Bank.
- Arndt C., Hussain M.A., Jones E.S., Nhate V., Tarp F. and Thurlow, J. (2012a). 'Explaining the Evolution of Poverty: The Case of Mozambique'. *American Journal of Agricultural Economics*, 94(4): 854-72.
- Arndt C., Distant R., Hussain M.A., Østerdal L.P., Huong P. and Ibraimo M. (2012b). 'Ordinal Welfare Comparisons with Multiple Discrete Indicators: A First Order Dominance Approach and Application to Child Poverty'. *World Development*, 40(11): 2290-301.
- Atkinson A., and Bourguignon F. (1982). 'The comparison of multi-dimensioned distributions of economic status'. *The Review of Economic Studies*, 49(2): 183-201.
- Atkinson A., and Lugo M.A. (2010). 'Growth, poverty and distribution in Tanzania'. Mimeo. International Growth Centre.
- Batana Y., and Duclos J. (2010). 'Comparing Multidimensional Poverty with Qualitative Indicators of Well-Being'. CIRPEE Working Paper 10-04.
- Bourguignon F., and Chakravarty S. (2003). 'The measurement of multidimensional poverty'. *Journal of Economic inequality*, 1(1): 25-49.
- Deaton, A. and V. Kozel. (2005) Data and dogma: The great Indian poverty debate. *World Bank Research Observer*. 20(2), 177-199.
- DNEAP (2010). *Poverty and Well-being in Mozambique: The third national assessment*. National Directorate of Studies and Policy Analysis, Ministry of Planning and Development, Maputo, Mozambique.
- DNPO (1998). *Poverty and Well-Being in Mozambique: The First National Assessment (1996-1997)*. Maputo and Washington DC: Ministry of Planning and Finance, Universidade Eduardo Mondlane, International Food Policy Research Institute.

- Duclos J., Sahn D. and Younger S. (2007). 'Robust multidimensional poverty comparisons with discrete indicators of well-being'. In Jenkins S.P. and Micklewright J. (eds), *Inequality and poverty reexamined*, Oxford: Oxford University Press.
- Elbers C., Lanjouw J. and Lanjouw P. (2003). 'Micro-level Estimation of Poverty and Inequality'. *Econometrica*, 71(1): 355-64.
- Elbers C., Lanjouw P. and Leite P. (2008). 'Brazil within Brazil: testing the poverty map methodology in Minas Gerais'. Washington DC: World Bank Policy Research Working Paper Series.
- Gordon D., Nandy S., Pantazis C., Pemberton S. and Townsend P. (2003). *Child poverty in the developing world*. Bristol: Policy Press.
- Grant S., Kajii A. and Polak B. (1992). 'Many good choice axioms: when can many-good lotteries be treated as money lotteries?'. *Journal of Economic Theory*, 56: 313-37.
- Grimm M., and Gunther I. (2007). 'Growth and poverty in Burkina Faso: A reassessment of the paradox'. *Journal of African Economies*, 16(1): 70-101.
- INE (2010) *Inquerito Sobre Orçamento Familiar, 2008/09: Quadros Basicos*. Maputo: Instituto Nacional de Estatística.
- Lehmann, E.L. (1955). 'Ordered families of distributions'. *The Annals of Mathematical Statistics*, 26(3): 399-419.
- Levhari, D., Paroush, J., and Peleg, B. (1975). 'Efficiency analysis for multivariate distributions'. *Review of Economic Studies*, 42(1): 87-91.
- Molina I., and Rao J. (2010). 'Small area estimation of poverty indicators'. *Canadian Journal of Statistics*, 38(3): 369-85.
- Müller, A., and Stoyan D. (2002). *Comparison Methods for Stochastic Models and Risks*. Chichester: John Wiley and Sons.
- National Bureau of Statistics (2009). *Household budget survey 2007*. Dar es Salaam: National Bureau of Statistics.
- National Bureau of Statistics (2010). *National Panel Survey Report, Round 1, 2008-2009*. Dar es Salaam: National Bureau of Statistics.
- National Bureau of Statistics (2012). *National Panel Survey Report, Wave 2, 2010-2011*. Dar es Salaam: National Bureau of Statistics.
- Preston, C. (1974). 'A generalization of the FKG inequalities'. *Communications in Mathematical Physics*, 36: 233-41.
- Shaked, M., and Shanthikumar, J.G. (2007). *Stochastic Orders*. New York: Springer.
- Strassen, V. (1965). 'The existence of probability measures with given marginals'. *Annals of Mathematical Statistics*, 36(2): 423-39.
- Tarozzi A. and Deaton A. (2009). 'Using census and survey data to estimate poverty and inequality for small areas'. *The Review of Economics and Statistics*, 91(4): 773-92.
- Østerdal, L.P. (2010). 'The mass transfer approach to multivariate discrete first order stochastic dominance: Direct proof and implications'. *Journal of Mathematical Economics*, 46: 1222-8.