

Poverty Mapping Based on First-Order Dominance with an Example from Mozambique

Arndt, Thomas Channing; Hussain, M. Azhar; Salvucci, Vincenzo; Tarp, Finn; Peter Østerdal, Lars

Published in:
Journal of International Development

DOI:
[10.1002/jid.3200](https://doi.org/10.1002/jid.3200)

Publication date:
2016

Document Version
Publisher's PDF, also known as Version of record

Citation for published version (APA):
Arndt, T. C., Hussain, M. A., Salvucci, V., Tarp, F., & Peter Østerdal, L. (2016). Poverty Mapping Based on First-Order Dominance with an Example from Mozambique. *Journal of International Development*, 28(1), 3-21.
<https://doi.org/10.1002/jid.3200>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact rucforsk@kb.dk providing details, and we will remove access to the work immediately and investigate your claim.

POVERTY MAPPING BASED ON FIRST-ORDER DOMINANCE WITH AN EXAMPLE FROM MOZAMBIQUE

CHANNING ARNDT¹, AZHAR M. HUSSAIN², VINCENZO SALVUCCI³, FINN TARP^{1,3*}
and LARS PETER ØSTERDAL⁴

¹UNU-WIDER, Helsinki, Finland

²Roskilde University, Roskilde, Denmark

³University of Copenhagen, Copenhagen, Denmark

⁴University of Southern Denmark, Odense, Denmark

Abstract: We explore a novel first-order dominance (FOD) approach to poverty mapping and compare its properties to small-area estimation. The FOD approach uses census data directly, is straightforward to implement, is multidimensional allowing for a broad conception of welfare and accounts rigorously for welfare distributions in both levels and trends. An application to Mozambique highlights the value of the approach, including its advantages in the monitoring and evaluation of public expenditures. We conclude that the FOD approach to poverty mapping constitutes a useful addition to the toolkit of policy analysts. © 2015 UNU-WIDER. Journal of International Development published by John Wiley & Sons, Ltd.

Keywords: first-order dominance; poverty mapping; small-area estimation; multidimensional poverty measurement; public expenditure evaluation; Mozambique

JEL classification: C15; C81; I32

1 INTRODUCTION

Information on welfare at small scales (such as the district or even village) is a valuable analytical input into decision-making about the allocation of government funds, regional planning and general policy formulation. The wide application of small-area estimation techniques (Elbers *et al.*, 2003; Molina & Rao, 2010), which estimate consumption poverty rates at small scales, attests to the value of the information generated. Small-area estimates underlie the subsequent elaboration of poverty maps, which allow a visual representation of the geographic dispersion of consumption poverty as well as (data permitting) changes in the geographic distribution of poverty levels through time.

*Correspondence to: Finn Tarp, UNU-WIDER, Helsinki, Finland.

E-mail: tarp@wider.unu.edu

This study has two main purposes. First, we argue that a multidimensional supplement to the current small-area methodology could add significant additional value, particularly in regions characterised by high levels of absolute poverty. Second, we present an analytical approach based on first-order dominance (FOD) that is a strong candidate to serve as that supplement.

Compared with small-area estimation, the FOD approach has the following desirable characteristics. First, it uses census data directly with a minimum of assumptions imposed. Second, the underlying concepts are simple, and the methodology is straightforward to implement. Third, the FOD approach is multidimensional, allowing for a broad conception of poverty. Fourth, FOD indicators can be chosen such that they relate directly to public expenditure priorities. Arguably, when a welfare measure for small areas is used to help guide the allocation of public expenditures on items such as water, sanitation, education and electrification, then direct indicators associated with these expenditure priorities should be relied upon.

The basic idea behind the proposed FOD approach is to construct poverty maps from multiple comparisons of districts (small areas). An FOD comparison determines if one district is unambiguously better off than another (for given indicators). By comparing each district with all others, it is possible to provide an intuitive ranking of all districts via the Copeland (1951) method. This approach is analogous to the way in which teams in a sports league are ranked by assigning points to wins, draws and losses from head-to-head matchups. For example, a winner gets two points, a draw gives one point to each team and a loss gives no points. The final ranking depends on the sum of points gained during the tournament.

The remainder of this article is structured as follows. To set the analytical scene, Sections 2 and 3 provide background on respectively small-area estimation and the FOD methodology, including an assessment of strengths and weaknesses. Sections 4 and 5 present our example application to Mozambique. Section 4 covers data and variables used, and Section 5 summarises results including a set of comparisons of traditional small-area and FOD estimates. Section 6 concludes that the FOD approach is a useful addition to the poverty measurement toolkit.¹

2 SMALL-AREA ESTIMATION

Small-area estimation techniques are applied in cases where there exists:

- (i) A sample survey with information on consumption (y) and household characteristics (X). As a sample will not cover all households in the targeted population, the ability to make viable inferences with respect to the welfare status of a relevant small subpopulation is determined by the sampling procedure.
- (ii) A census of the population that occurred reasonably close in time to the household consumption survey. Censuses, as a rule, do not attempt to obtain consumption information from all households. Instead, they typically aim to obtain information on household characteristics that are relatively easily observable for all, or a large

¹The idea of applying multidimensional techniques to rank welfare in small areas has been advanced by Arndt *et al.* (2012b), in a UNU-WIDER working paper version of this study by Arndt *et al.* (2013), and by Permanyer (2013), who introduces human development index-oriented methods to explore the spatial distribution of welfare. Permanyer *et al.* (2014) apply these techniques to census data from 13 African countries. See also Sonne-Schmidt *et al.* (2008, 2015) for background.

subsample of, households in the target population. If the census and survey are designed with small-area estimation in mind, a set of collected household characteristics (X) will be comparable between the census and the survey.

When these two elements are present, small-area estimation can proceed by relying on a set of domain-specific survey-based regressions that model (per capita) log consumption, y , as a function of explanatory household-level and area-level variables, X^{sur} , derived from the household survey.

In this way, a vector of estimated parameters $\hat{\beta}^{sur}$ emerges where superscript *sur* indicates that the variable/parameter is from the survey. These parameters are in turn combined with explanatory household-level and area-level variables from the census (X^{cen}), making it possible for the analyst to assign an expected household (log) consumption level $\hat{y} = \hat{\beta}^{sur} X^{cen}$ to each household in the census along with its estimated variance. Based on this information, poverty measures, such as the headcount, can be derived for households in a given small area (Foster *et al.*, 1984).²

Such small-area estimation exercises have provided welcome and valuable information across a wide range of countries. However, they also have some disadvantages.

- (i) It is sometimes not possible to implement the small-area estimation approach and this is so 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 through joint planning so that the statistical authority conducts household consumption surveys close to the next census in a coordinated fashion.
- (ii) The small-area methodology depends on the measurement of nominal consumption in the survey as well as the estimation of poverty lines, which should reflect a (reasonably) constant living standard across space and time. This remains challenging and is frequently controversial both with respect to national numbers and regional poverty profiles.³
- (iii) Poverty estimates in household surveys will tend to 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 percentage point decline in poverty they observed in a 5-year period in Burkina Faso to drought conditions experienced in the initial period. 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 2002/2003 to 2008/2009 and a 12 percentage point increase in poverty in the rural areas of Sofala and Zambezia provinces over the same period. Accordingly, changes in measured consumption poverty across time and space represent a difficult-to-decompose mix of underlying development progress/regress/stagnation and ephemeral shocks to welfare that may be positive or negative, in combination with sample and non-sample error.

²In principle, the small-area methodology can be used to develop any indicator simply by varying the left-hand side variable.

³See, for example, Atkinson and Lugo (2010); Alfani *et al.* (2012); Grimm and Gunther (2007); Deaton and Kozel (2005).

- (iv) It is assumed 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 subdomain spatial correlation.⁴ In consumption regressions, unobserved heterogeneity across survey strata (e.g. across provinces) is often absorbed through dummy variables to capture domain specificity. These average characteristics of the 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, underlining the potential complexity of the consumption regressions.

We finally recall that, as these estimates are based on consumption, they provide a one-dimensional measure of poverty. Therefore they do not *per se* respond to Sen (1985) who argues that poverty must be treated as a multidimensional phenomenon. In light of these issues, there appears to be scope for an expanded toolkit that can help complement small-area estimates and provide rigorous welfare rankings across space and through time based on census data. We turn now to a candidate approach to include in this expanded toolkit.

3 FIRST-ORDER DOMINANCE METHODOLOGY

Existing literature on 'robust' methods for comparing multidimensional population welfare, poverty and inequality (e.g. Atkinson & Bourguignon, 1982; Duclos *et al.*, 2007; Bourguignon & Chakravarty, 2003; Batana & Duclos, 2010) relies on stochastic dominance concepts for comparisons that are valid for broad classes of underlying social welfare functions. The FOD criterion, in specific, corresponds to what in probability theory is referred to as the *usual (stochastic) order* (Lehmann, 1955).⁵ This implies that the FOD approach does not depend on a weighting scheme or on strongly simplifying assumptions about the second order and cross derivatives of the social welfare function (Arndt *et al.*, 2012b). Instead, for the case of binary welfare indicators where individuals or households are either deprived or not deprived in each specific welfare dimension, the FOD criterion simply asserts that it is better to be not 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 $2^5 = 32$ different possible states corresponding to whether they are deprived or not deprived in the various dimensions.⁶ 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.

⁴This is potentially important because Tarozzi and Deaton (2009) note that small violations of the 'area homogeneity' assumption may result in misleading inference. They also express concerns that subdomain spatial integration might not be completely taken into account, which leads to an underestimation of sample errors of small area estimates. In an evaluation of the poverty mapping methodology conducted on Brazilian data, both of these concerns seemed to be of minor importance (Elbers *et al.*, 2008).

⁵For a general treatment of stochastic dominance theory, we refer to Müller and Stoyan (2002) or Shaked and Shanthikumar (2007).

⁶Note that the curse of dimensionality will limit the number of indicators considered even for very large data sets.

Intermediate rankings are somewhat more complex. 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 is always at least equivalent to 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 information/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 judgement is made as to the relative importance of dimension 3 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.⁷ Helpfully, this condition can be defined as a network flow problem with limitations on allowed transfer paths (e.g. Preston, 1974; Mosler & Scarsini, 1991).

In the applications considered in the succeeding text, 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 cases. Define binary indices i, j, k, l and m , which each can take the value 0 or 1. As discussed in Section 4.2 in the succeeding text, i, j, k, l and m may be for example water, sanitation, electricity, education and durable goods. As in the previous text, the value 1 refers to not deprived and the value 0 to deprived for the five dimensions. Define binary indices i', j', k', l' and 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, an allowable transfer is where $i' \leq i, j' \leq j, k' \leq k, l' \leq l$ and $m' \leq m$. Transfer of probability mass from state $(ijklm)$ to state $(ijklm)$ is also excluded for numerical reasons. Under these conditions, population A first-order dominates population B if and only if the following linear programme is feasible.

$$\text{Min } y = 1 \quad (1)$$

subject to

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

With the ability to compare any two populations, large numbers of subpopulation 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 $100^2 - 100 = 9900$ comparisons. Following Copeland (1951),

⁷This is equivalent to the condition that population A has higher welfare than population B for *any* increasing social welfare function; compare with Strassen (1965), Levhari *et al.* (1975), and Grant *et al.* (1992) (see also Østerdal, 2010).

complete welfare rankings of regions can be generated by, for example, counting the number of times a given region dominates other regions and subtracting the number of times the same region is dominated by other regions generating a score in the interval $[-99, 99]$. Regions can then be naturally ranked with higher scores superior to lower scores (for the properties of this approach, see Saari and Merlin, 1996, and Merlin & Saari, 1997). We then define a 'Copeland index' whereby all scores are normalised to fall in the interval $[-1, 1]$.

Note that transitivity applies meaning that if region A first-order dominates region B and region B first-order dominates region C, then region A must first-order dominate region C.⁸ Therefore, if A first-order dominates B, then the Copeland index score of A must be strictly greater than the score of B (unless A and B are equivalent).

There are two main challenges in applications of the FOD approach:

- (i) FOD comparisons can be indeterminate. If neither A dominates B nor B dominates A (e.g. both linear programmes are infeasible), then the welfare rankings of A and B are indeterminate without further information.
- (ii) While A may dominate B, the degree of domination is unknown without further information. Dominance may therefore reflect both fine differences in the distributions across states or may reflect substantially different circumstances between A and B.

While recognising these challenges, it is also pertinent to note that the multiple comparisons inherent in the evaluation of numerous small areas across space generate additional information that helps offset them. Suppose that neither A nor B dominates the other, but that on net A dominates 20 other regions while B dominates negative one (i.e. 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 as in the Copeland index.

Because many countries have censuses from multiple years, it is also possible to use the FOD criterion to determine whether welfare has been improving through time at various geographical scales. The comparison of interest, region D at time t versus region D at time $t+s$, naturally yields only one comparison pair. However, use of bootstrapping can help to mitigate the two disadvantages associated with the FOD approach through the generation of multiple comparisons.

Bootstrapping can also be applied to the spatial case. While the marginal gain in information generated by bootstrapping spatial comparisons across small areas is perhaps not as important as in the temporal case (due to the already large numbers of comparisons normally available when census data are employed), the potential information gain can be obtained at relatively low cost. Some additional computer time and data management are all that is required.⁹

⁸This 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.

⁹Finally, while bootstrapping census data is not normally performed, the principles and concepts remain straightforward. Indeed, bootstrapping the census is simply random sampling from a population with replacement.

Before moving on to empirical applications, two additional observations are pertinent. First, while large numbers of comparisons generate substantial additional information, the limitations associated with applying a strict criterion, such as FOD, cannot be avoided entirely.¹⁰ Consider a simplified example where the FOD criterion is applied to individuals in a population based on five welfare indicators. Suppose individual A in this population is not deprived in dimension 1 and is deprived in all other dimensions [i.e. (1, 0, 0, 0, 0)]. Further, suppose that all other individuals are characterised as (0, j , k , l , m) where j , k , l and m are binary variables taking values zero or one and $j+k+l+m \geq 1$. In this instance, individual A neither dominates nor is dominated by any other individual in the population. The net domination score for individual A is zero. More generally and returning to comparisons across populations, for the case where the welfare profiles differ dramatically, the FOD criteria may yield relatively little insight.

On the other hand, this limitation does not apply through time. Failure to advance through time implies that the distributional changes observed over time do not represent an unequivocal improvement over conditions that existed in the past. Here, the strict comparative nature of the FOD criterion becomes an advantage. Specifically, the FOD approach requires (non-negative) progress across all indicators *and* across the range of the welfare distribution. In addition, progress for the poorest is required.

4 MOZAMBIQUE APPLICATION: DATA AND EMPIRICAL CHOICES

The data sets used here are the 1997 and 2007 censuses, and household consumption surveys from 1996/1997 (*Inquérito aos Agregados Familiares 1996/1997*) and 2008/2009 (*Inquérito sobre Orçamento Familiar 2008/2009*) from Mozambique. The surveys and censuses were conducted by the National Statistical Institute. The 1997 and 2007 censuses surveyed the full population, which amounted to about 15 and 21 million people, respectively. The surveys are representative for the whole of Mozambique as well as for the rural and urban zones and each of the 10 provinces plus Maputo City. The sample in 1996/1997 consists of 8274 households, whereas 10832 households were interviewed in 2008/2009. Both surveys include information on general characteristics of individuals and households, and there is information on daily, monthly and own consumption, as well as information on possession of durable goods, transfers and gifts. These surveys have been used to estimate both the first and also the latest set of official poverty rates at the national and regional levels. The headcount ratios presented in the following analyses are computed using region-specific poverty lines that explicitly consider provincial as well as rural/urban differences (DNPO, 1998; INE, 2010; DNEAP, 2010). In general, these poverty lines are lower than the standard \$1.25 measure.

4.1 Variables Used for Small-Area Estimation

The information included in the small-area analysis presented here is limited by what is available in both the census and in the survey. The common information available covers

¹⁰In addition, the Copeland method has certain limitations, as discussed in Saari and Merlin (1996) and Merlin and Saari (1997).

demographic characteristics, education, assets, own production of food items and labour market variables. The same set of candidate variables is 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,¹¹ 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.¹² Details on the correction procedure can be found in DNEAP (2010).

4.2 First-Order Dominance Welfare Indicators

Just as traditional small-area estimates are dependent on quality household consumption information, including price information to ascertain differences in costs of living, an FOD-based metric of relative welfare depends on the quality and relevance of its constituent indicators. Available indicators must address meaningful welfare dimensions that can be consistently interpreted within relevant domains of space and time.¹³ The choice of indicators will influence results. Accordingly, the FOD approach does not obviate, in any way, the need for careful consideration of the welfare indicators sought in census-type questionnaires.

For the case of Mozambique, five welfare indicators are considered here—inspired by the notion of ‘severe deprivation’ based on the Bristol indicators (Gordon *et al.*, 2003).¹⁴ These indicators represent a reasonable set of binary indicators given data availability and prevailing living conditions. For 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 available in 2007, the definitions differ slightly between 1997 and 2007 for the spatial analyses (within-year comparisons). For the temporal analysis, the coarser definitions from 1997 must also be applied in 2007.

The indicators are as follows:

- *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.

¹¹The dependency ratio is the ratio of dependents to the working-age population. Dependents are usually defined as people younger than 15 years or older than 64, while the working-age population is people aged 15–64 years.

¹²Evidence from Mozambique and other countries suggests that, particularly in urban areas, people tend to underreport consumption in part due to a more diversified diet and higher food consumption outside the home. This probable non-sampling error in self-reported food consumption can affect estimated poverty levels. For Mozambique, this overestimates poverty rates by about 3 percentage points nationally.

¹³Correlations across indicators function similarly to multicollinearity in linear regression. They reduce the effective quantity of information that is brought to bear. This is most easily seen in the case of two perfectly correlated indicators. In this case, the analysis is effectively taking place in one fewer dimension.

¹⁴Gordon *et al.* (2003) define severe deprivation as failure to attain even most basic thresholds for a series of common wellbeing indicators. The deprivation thresholds used to define severe deprivation are then lower than those frequently published by international organisations. For example, access only to surface water (e.g. rivers and ponds) for drinking or no access to a toilet of any kind.

- *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 who have at least one household member with some education.¹⁵
- *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 (per cent)

| 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 |

N = 146 (districts). Source: Authors' analysis based on the 1997 and 2007 censuses, and (only for the non-poverty indicator) the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

4.3 First-Order Dominance Metric—The Spatial Case

There are 146 districts in Mozambique identified in both 1997 and 2007, giving a total of 21 170 comparisons for each year without the bootstrap. To facilitate comparing these FOD spatial analyses for 1997 and 2007 with the traditional small-area poverty rate estimates, we define a measure of dominance labelled *spatial FOD index*, which is a 0–1 renormalized version of the Copeland index discussed in Section 3 with higher values corresponding to lower rankings (like the poverty rate).¹⁶

The bootstrap approach referred to in Section 3 is also employed in the analysis. The size of each bootstrap sample is equal to the number of households in the least populous district, 1828 households. One-hundred bootstrap repetitions are employed generating more than two million spatial comparisons for each census year and 200 temporal comparisons for each district (see next section). Because of the reasonably large number of spatial comparisons conducted without the bootstrap, net domination measures generated with and without the bootstrap give very similar results. We therefore opt to present the spatial results without the bootstrap.

¹⁵This indicator is relevant for Mozambique, where educational attainment for the period in question was very low (Table 1). Existing household budget surveys indicate that households with at least some education are better off (DNPO, 1998; DNEAP, 2010).

¹⁶The affine transformation is $S = -(C - 1)/2$ where S is spatial FOD index and C is the Copeland index.

4.4 First-Order Dominance Metric—The 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 vice versa. For each district, we define three possible results:

0: neither 2007 dominates 1997 nor 1997 dominates 2007

1: 2007 dominates 1997

−1: 1997 dominates 2007.

We apply the bootstrap to the temporal case in order to generate probabilistic measures of dominance. In the event, 1997 never dominates 2007 for any bootstrap draw across all districts. 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 temporal domination called the *temporal FOD index*.

5 RESULTS

Figure 1 illustrates the small-area poverty estimation results for 1996/1997 (panel a), for 2008/2009 (panel b) and for the change in the headcount ratio between 1996/1997 and 2008/2009 (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 small-area estimation poverty map in Figure 1 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 in 2007. Instead, the temporal FOD index provides a measure of change through time. Descriptive statistics for the estimated indices are displayed in Table 2.

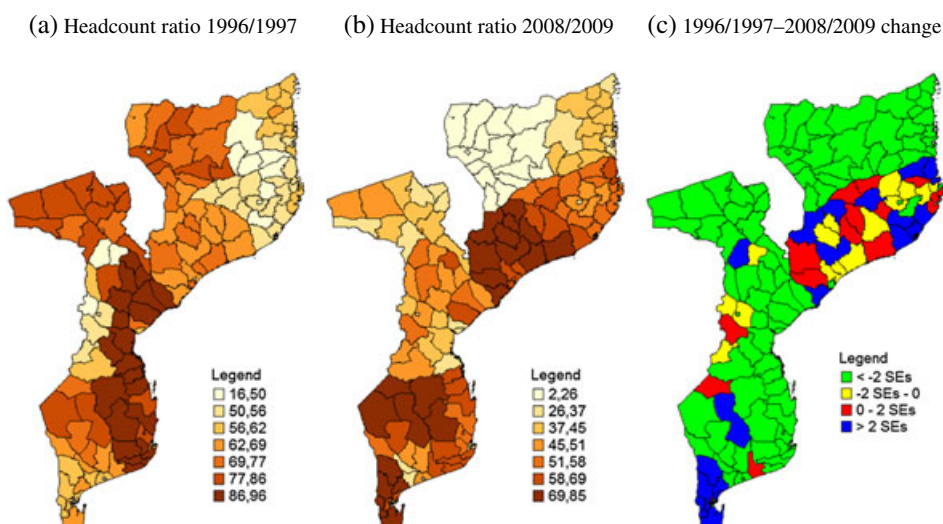


Figure 1. Small-area district poverty estimates, 1996/1997 and 2008/2009 (per cent). Source: Authors' analysis based on the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

(a) Spatial FOD index 1997 (b) Spatial FOD index 2007 (c) Temporal FOD index

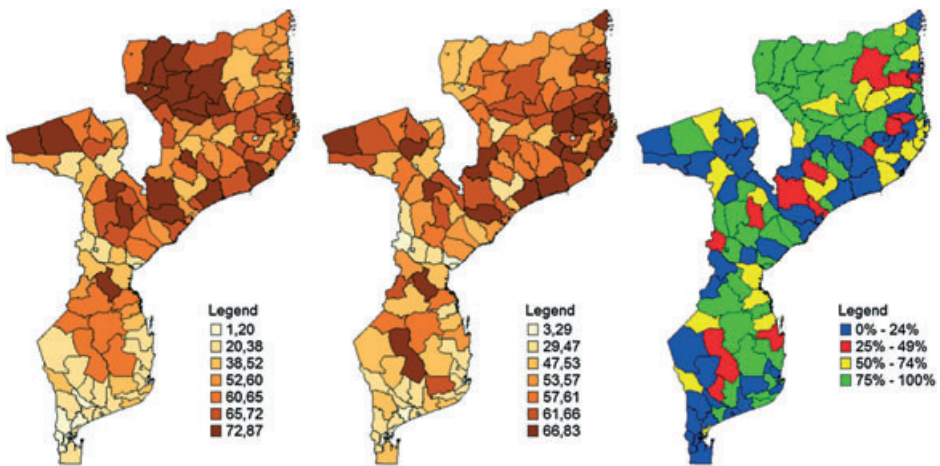


Figure 2. First-order dominance (FOD) small-area district mapping, 1997 and 2007. In panels a and b, lower numbers represent superior district rankings while in panel c, higher numbers represent higher probability of progress. Source: Authors' analysis based on the 1997 and 2007 censuses from Mozambique.

Table 2. Descriptive statistics for the indices used

| Index and year | Mean | St. dev. | Median | Min | Max |
|---------------------------------------|-------|----------|--------|-------|-------|
| Headcount ratio 1996/1997 | 66.9 | 15.0 | 65.1 | 16.4 | 96.0 |
| Headcount ratio 2008/2009 | 47.5 | 16.9 | 49.0 | 2.3 | 84.8 |
| Headcount change, 1996/1997–2008/2009 | –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 |

N = 146 (districts). Source: Authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

5.1 Small-Area Estimation 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.¹⁷ 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 2008/2009, the small-area poverty map changes. Of the districts placed in the poorest of the seven categories in 1997, only two districts remain among the poorest. Districts

¹⁷In 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 1996/1997 for the 2008/2009 poverty map.

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 2008/2009, 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 district poverty trends from 1996/1997 to 2008/2009 is shown in panel c of Figure 1. Here, the districts in which consumption poverty decreased by more than two standard errors 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 by less than two standard errors are marked in red, while districts in which it increased by more than two standard errors are in blue.¹⁸ Looking at the district-level poverty change, 77 per cent of the districts (accounting for 70 per cent of the entire population) experienced consumption poverty reduction over the decade from 1996/1997 to 2008/2009. 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 (30 per cent of the population) 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 1996/1997 saw the greatest reductions in poverty rates. In this sense, the reduction in consumption poverty over 1996/1997 to 2008/2009 can be said to have a pro-poor bias.

5.2 First-Order Dominance Results

The 146 districts in Figure 2 showing the FOD results are coloured as well depending on seven ordered levels for the spatial FOD index. Also in this case, the levels are chosen such that (roughly) an equal number of districts are in each level. In Figure 2 panel a (corresponding to 1997), we see that the districts ranked as most deprived are those located in the northern and central areas. The share of the population living in the worst performing districts—the bottom three categories—is about 33 per cent. None of the southern districts appear in the worst-ranked group. As it emerged in the headcount ratio analysis, the FOD results 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 lowest-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. In this case, a slightly higher share of the population—40 per cent—was located in the three worst performing district categories.

Looking at the FOD temporal index, no district in 2007 is, as already noted, dominated by itself in 1997. There are 76 districts (out of 146, corresponding to 45 per cent of the

¹⁸Note that the standard error is calculated about the mean while the categories are defined about zero.

population) 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 fairly broad-based advance in living conditions between 1997 and 2007.

Table 3 provides a useful summary of the district FOD results by province. The table shows the average rank of the districts within each province for each year, as well as the average probability of advance. For comparison purposes, the average headcount ratio rank and the average headcount ratio change in rank are displayed. Note that only three provinces gain in average FOD ranking: Niassa, Manica and Sofala. These provinces are associated with strong probabilities of advance. Two provinces exhibit probabilities of advance greater than 50 per cent but declines in average ranking. These are Inhambane, which shows only a minor decline in average ranking, and Cabo Delgado. The latter province shows the third-highest probability of gain but loses in terms of average ranking. This is explained in large part by the movement of districts in Niassa from an average ranking substantially below Cabo Delgado in 1997 to an average ranking materially above Cabo Delgado in 2007.

As noted, Maputo Province and Maputo City ranked well ahead of other provinces in 1997 and they retain their leadership rankings despite the weakest performance by far in terms of average probability of gain. The remaining provinces exhibit probabilities of gain of about 33 per cent and mild declines in rankings. Overall, there is solid coherence between the spatial rankings and the temporal gain probabilities. Finally, sensitivity analysis (not shown) does not reveal any district where the FOD procedure exhibits particularly weak power (indeterminacy) in spatial comparisons.

Table 3. Average district spatial FOD ranking, 1997 and 2007, average probability of 2007 FOD 1997, average district headcount ratio rank, 1997 and 2007, and average change in the headcount ratio 1997–2007, by province

| Province | Average district FOD rank 1997 | Average district FOD rank 2007 | Average probability of 2007 FOD 1997 | Average district headcount ratio rank 1997 | Average district headcount ratio rank 2007 | Average change in headcount ratio 1997–2007 |
|--------------|--------------------------------|--------------------------------|--------------------------------------|--------------------------------------------|--------------------------------------------|---------------------------------------------|
| Niassa | 106.8 | 68.1 | 90.7 | 83.5 | 12.1 | 46.4 |
| Cabo Delgado | 70.6 | 85.5 | 62.4 | 30.0 | 44.7 | 15.6 |
| Nampula | 84.3 | 96.8 | 37.8 | 30.0 | 84.3 | 1.2 |
| Zambezia | 93.4 | 100.7 | 27.6 | 78.6 | 127.8 | –1.7 |
| Tete | 67.5 | 79 | 35.5 | 113.3 | 44.8 | 43.3 |
| Manica | 51.8 | 44.6 | 57.9 | 26.1 | 68.4 | 5.7 |
| Sofala | 59.8 | 50.9 | 76.8 | 105.8 | 53.1 | 40.9 |
| Inhambane | 39.4 | 41.1 | 57.4 | 121.9 | 96.5 | 28.4 |
| Gaza | 26.6 | 30.8 | 31.5 | 62.7 | 67.7 | 16.7 |
| Maputo Prov. | 12.7 | 16.4 | 5.1 | 28.9 | 96.4 | –7.0 |
| Maputo City | 3.8 | 7.4 | 0 | 12.9 | 21.3 | 19.9 |

N = 146 (districts). District populations used as weights. Source: Authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

5.3 Comparing Small-Area and FOD Mapping

Comparing the traditional small-area estimates and the FOD small-area results for 1996–1997/1997, we find that the rankings of poorer and richer districts are visibly different: 47 districts change their ranking by more than 50 positions. In particular, most centre-northern districts are much better ranked in the small-area estimates than in the FOD mapping, while the opposite holds for the majority of districts in the centre-south. Similar results prevail when comparing the small-area mapping for 2008/2009 with the FOD mapping in 2007. In 2007, the two northern-most provinces score better in the small-area estimates than in the FOD analysis, while the three southern provinces (excluding Maputo City) are much better ranked in the multidimensional FOD.

We undertake a correlation analysis based on the rankings obtained so as to provide a finer overview of the differences. In Figure 3, scatter plots of the 146 Mozambican districts for 1996–1997/1997 (panel a) and 2008–2009/2007 (panel b) are displayed. On the horizontal axis, we show the district headcount ratio (1996/1997 and 2008/2009), and on the vertical axis, we show the corresponding spatial FOD index (1997 and 2007). The correlation coefficient between the

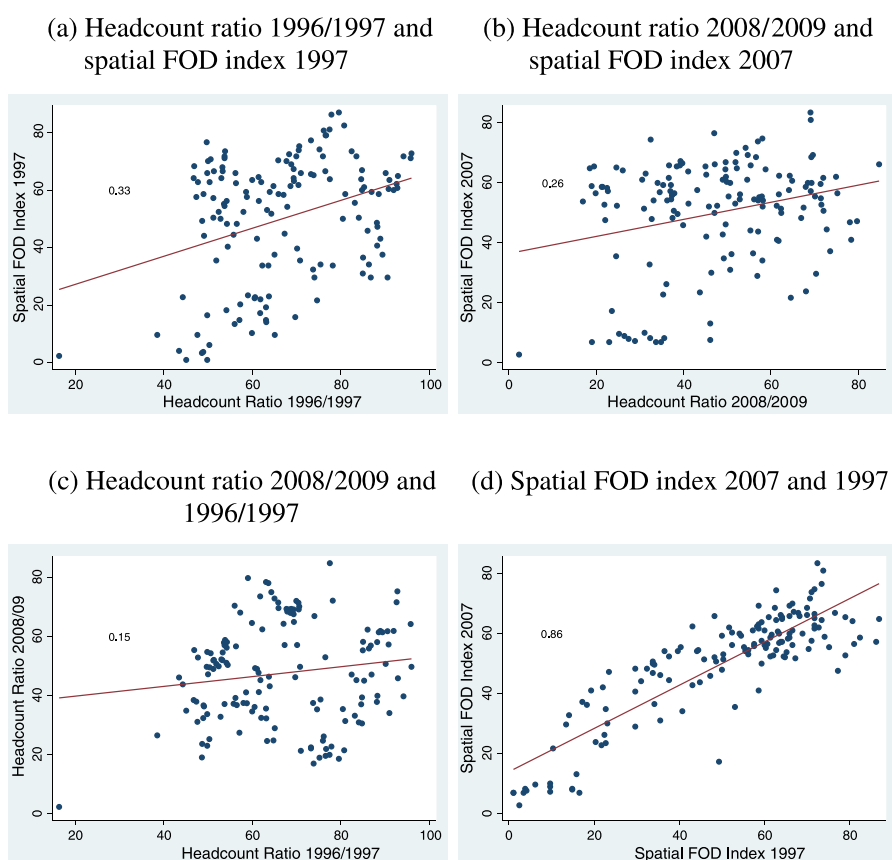


Figure 3. Correlation between headcount and first-order dominance (FOD) measures. The correlation coefficient is presented on the left-hand side of each figure. Source: Authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

two indices is 0.33 for 1996–1997/1997 and 0.26 for 2008–2009/2007. Concerning the changes over the analysed decade (not shown in the graph), a correlation of 0.33 is observed between the two welfare indices. This suggests 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 small-area poverty measure, effectively based on consumption, is strongly influenced by food availability and price (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 errors also undoubtedly contribute to the re-rankings.

In contrast, the five indicators underlying the FOD indices tend to be substantially 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 prone to non-sample error.

Volatility in the consumption measure and relative stability of the FOD indices are reflected in the correlations between the headcount from 1996/1997 and 2008/2009 and between the spatial FOD indices for 1997 and 2007. Indeed, the correlation coefficient between the headcount ratio in 1996/1997 and in 2008/2009 is low at 0.15 (Figure 3, panel c). Conversely, the correlation coefficient for the spatial FOD index in 1997 and in 2007 is 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 (i.e. provincial dummy variables) 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 provides 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 provincial urban and rural areas, to estimate welfare levels in all districts within that larger area. For the FOD, provincial boundaries are irrelevant to the calculations. Correspondingly, the implications of provincial boundaries are a lot less evident in Figure 2.¹⁹

5.4 Consumption Poverty as Indicator Instead of Radio in FOD Mapping

We next proceed to employ the FOD as an extension of the small-area methodology by substituting the poor/non-poor indicator from the small-area poverty analysis for the radio indicator among the five FOD variables. Not surprisingly, when consumption poverty is included, the FOD mapping becomes more similar to the results derived from the small-area poverty 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

¹⁹Hentschel *et al.* (2000) compare small area poverty estimates and indicator-based small-area deprivation and find greater consistency for Ecuador. Ecuador is, however, far wealthier than Mozambique with the likelihood of greater stability in consumption and greater correlations with other indicators.

(a) Spatial FOD index 1997

(b) Spatial FOD index 2007

(c) Temporal FOD index

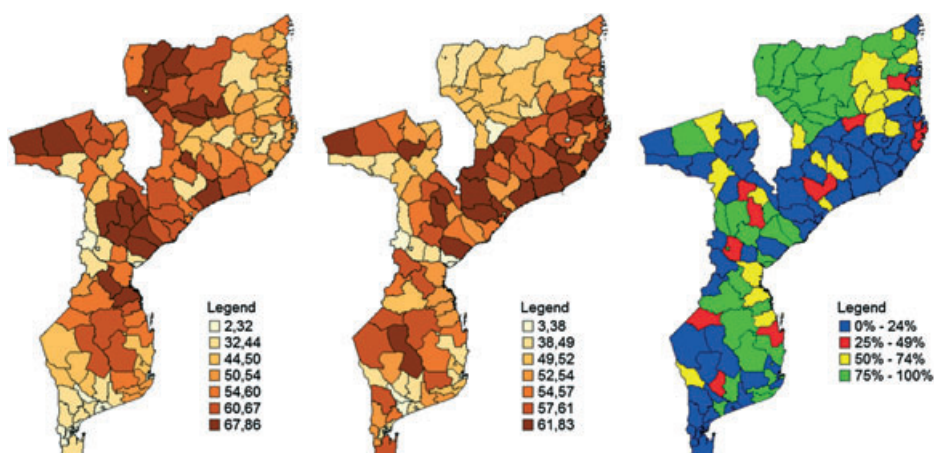


Figure 4. First-order dominance (FOD) mapping with non-poverty from poverty mapping as a welfare indicator, 1997 and 2007. Source: Authors' analysis based on the 1997 and 2007 censuses and the consumption surveys IAF 1996/1997 and IOF 2008/2009 from Mozambique.

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.

6 CONCLUSIONS

The FOD approach to small-area estimation represents a useful addition to the welfare analysis toolkit. The approach is flexible, straightforward to apply and multidimensional. Because of the large number of comparisons inherent in small-area estimation using census data, FOD appears to be well suited to the problem of devising welfare rankings for small areas. In addition, FOD provides a rigorous approach to evaluating progress through time at smaller or larger areas of disaggregation. In spatial comparisons, FOD metrics penalise unequal distributions of outcomes across areas. In temporal comparisons, FOD metrics penalise any backsliding in the distribution of outcomes through time. These features of the FOD approach are appealing and reflect the calls for such features by Permanyer *et al.* (2014) and the broader discussions on the 2030 Sustainable Development Goals approved by the UN General Assembly in September 2015.

For the example case of Mozambique, the FOD approach confirms broad-based progress between 1997 and 2007 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.

First-order dominance rankings of districts are relatively stable through time, especially when compared with traditional small-area poverty estimation, due, at least in part, to their basis in a relatively stable and easy to observe set of indicators. Differences with traditional small-area estimation stem from a variety of sources including (i) the metric

(e.g. consumption versus non-monetary measures), (ii) unidimensional versus multidimensional comparisons, (iii) the aggregation approach (e.g. FOD versus regression-based prediction) and (iv) different data sources (e.g. census versus combined survey/census) with different degrees of sample and non-sample error.

While pinpointing the exact sources of differences is effectively impossible, the FOD results are informative regardless of the source of difference. Consider each of the four potential sources of difference. Four of the five indicators employed to develop the FOD rankings relate directly to priority public expenditures in water, sanitation, education and infrastructure (electricity) (source i). Because poverty is widely recognised as a multidimensional phenomenon, multidimensional analysis is attractive (source ii). Sensitivity analysis furthermore reveals that the FOD results correlate highly with the multidimensional poverty index of Alkire and Foster (2011) and the Oxford Poverty and Human Development Initiative (2014).²⁰ This implies robustness in the differentials with small-area mapping (source iii).²¹ Finally, the direct use of multiple indicators from census data constitutes in our assessment an advantage (source iv).

Consequently, and particularly for the purposes of evaluating and prioritising government expenditures at local levels in low-income countries, the exact source of differential between the two approaches is not particularly germane. The FOD is likely to be informative whatever the weights on the alternative sources of difference might be. Overall, we believe that, properly conducted, the FOD approach constitutes a useful addition to the analytical toolkit whose scores/rankings provide additional insights of relevance to the policy-making process.

With respect to future research, the FOD approach allows robust welfare comparison 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 datasets, such as the Afrobarometer, may be one useful approach in moving forward.

REFERENCES

- Alfani F, Azzarri C, d'Errico M, Molini V. 2012. Poverty in Mozambique. New evidence from recent household surveys. Policy Research Working Paper 6217. World Bank.
- Alkire S, Foster J. 2011. Counting and multidimensional poverty measurement. *Journal of Public Economics* **95**: 476–487.
- Arndt C, Hussain MA, Jones ES, Nhate V, Tarp F, Thurlow J. 2012a. Explaining the evolution of poverty: the case of Mozambique. *American Journal of Agricultural Economics* **94**: 854–872.
- Arndt C, Distant R, Hussain MA, Østerdal LP, Huong P, Ibraimo M. 2012b. Ordinal welfare comparisons with multiple discrete indicators: a first order dominance approach and application to child poverty. *World Development* **40**: 2290–2301.
- Arndt C, Hussain MA, Salvucci V, Tarp F, Østerdal LP. 2013. Advancing Small Area Estimation. WIDER Working Paper 2013/53.
- Atkinson A, Bourguignon F. 1982. The comparison of multi-dimensioned distributions of economic status. *The Review of Economic Studies* **49**: 183–201.

²⁰The details of the sensitivity analysis are available upon request, and this correspondence reinforces the relevance of the FOD approach.

²¹See also Hussain (2014).

- Atkinson A, Lugo MA. 2010. Growth, poverty and distribution in Tanzania. International Growth Centre Working Paper 0831. International Growth Centre.
- Batana Y, Duclos J. 2010. Comparing multidimensional poverty with qualitative indicators of well-being. CIRPEE Working Paper 10–04. CIRPEE.
- Bourguignon F, Chakravarty S. 2003. The measurement of multidimensional poverty. *Journal of Economic Inequality* **1**: 25–49.
- Copeland AH. 1951. A 'reasonable' social welfare function. Seminar on Application of Mathematics to the Social Sciences. University of Michigan: Mimeo.
- Deaton A, Kozel V. 2005. Data and dogma: the great Indian poverty debate. *World Bank Research Observer* **20**: 177–99.
- DNEAP. 2010. Poverty and well-being in Mozambique: the third national assessment. In DNEAP (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). DNPO (National Directorate of Policy and Budget), Ministry of Planning and Finance, Universidade Eduardo Mondlane, International Food Policy Research Institute. Maputo, Mozambique.
- Duclos J, Sahn D, Younger S. 2007. Robust multidimensional poverty comparisons with discrete indicators of well-being. In *Inequality and Poverty Reexamined*, Jenkins SP, Micklewright J (eds). Oxford University Press: Oxford, UK.
- Elbers C, Lanjouw J, Lanjouw P. 2003. Micro-level estimation of poverty and inequality. *Econometrica* **71**: 355–364.
- Elbers C, Lanjouw P, Leite P. 2008. Brazil within Brazil: testing the poverty map methodology in Minas Gerais. Policy Research Working Paper Series. World Bank.
- Foster J, Greer J, Thorbecke E. 1984. A class of decomposable poverty measures. *Econometrica* **52**: 761–776.
- Gordon D, Nandy S, Pantazis C, Pemberton S, Townsend P. 2003. Child Poverty in the Developing World. Policy Press: Bristol, UK.
- Grant S, Kajji A, Polak B. 1992. Many good choice axioms: when can many-good lotteries be treated as money lotteries? *Journal of Economic Theory* **56**: 313–337.
- Grimm M, Gunther I. 2007. Growth and poverty in Burkina Faso: a reassessment of the paradox. *Journal of African Economies* **16**: 70–101.
- Hentschel J, Lanjouw J, Lanjouw P, Poggi J. 2000. Combining census and survey data to trace the spatial dimensions of poverty: a case study of Ecuador. *The World Bank Economic Review* **14**: 147–165.
- Hussain MA. 2014. EU country rankings' sensitivity to the choice of welfare indicators. *Social Indicators Research* forthcoming. DOI:10.1007/s11205-014-0830-4.
- INE. 2010. Inquerito Sobre Orçamento Familiar, 2008/2009—Quadros Basicos. INE (Instituto Nacional de Estatística): Maputo, Mozambique.
- Lehmann EL. 1955. Ordered families of distributions. *The Annals of Mathematical Statistics* **26**: 399–419.
- Levhari D, Paroush J, Peleg B. 1975. Efficiency analysis for multivariate distributions. *Review of Economic Studies* **42**: 87–91.
- Merlin VR, Saari DG. 1997. Copeland method II: manipulation, monotonicity, and paradoxes. *Journal of Economic Theory* **72**: 148–172.
- Molina I, Rao J. 2010. Small area estimation of poverty indicators. *Canadian Journal of Statistics* **38**: 369–385.
- Mosler KC, Scarsini M. 1991. Some theory of stochastic dominance. In *Stochastic Orders and Decision Under Risk*, Mosler KC, Scarsini M (eds). Institute of Mathematical Statistics: Hayward, California.

- Müller A, Stoyan D. 2002. Comparison Methods for Stochastic Models and Risks. John Wiley & Sons: New York.
- Oxford Poverty and Human Development Initiative. 2014. Mozambique country briefing. In Multidimensional Poverty Index Data Bank. OPHI (Oxford Poverty and Human Development Initiative), University of Oxford: Oxford, UK.
- Permanyer I. 2013. Using census data to explore the spatial distribution of human development. *World Development* **46**: 1–13.
- Permanyer I, Esteve-Palos A, Garcia J, McCaa R. 2014. Human development index-like small area estimates for Africa computed from IPUMS-international integrated census microdata. *Journal of Human Development and Capabilities: A Multi-Disciplinary Journal for People-Centered Development*. DOI:10.1080/19452829.2014.956300.
- Preston C. 1974. A generalization of the FKG inequalities. *Communications in Mathematical Physics* **36**: 233–41.
- Saari DG, Merlin VR. 1996. The Copeland method: I.: relationships and the dictionary. *Economic Theory* **8**: 51–76.
- Sen AK. 1985. Commodities and Capabilities. Oxford University Press: Oxford.
- Shaked M, Shanthikumar JG. 2007. Stochastic Orders. Springer Science and Business Media: New York.
- Sonne-Schmidt C, Tarp F, Østerdal LP. 2008. Ordinal Comparison of Multidimensional Deprivation: Theory and Application. Discussion Paper 08-33 Department of Economics, University of Copenhagen.
- Sonne-Schmidt C, Tarp F, Østerdal LP. 2015. Ordinal Bivariate Inequality: Concepts and Application to Child Deprivation in Mozambique. *Review of Income and Wealth*. DOI:10.1111/roiw.12183.
- Strassen V. 1965. The existence of probability measures with given marginals. *Annals of Mathematical Statistics* **36**: 423–439.
- Tarozzi A, Deaton A. 2009. Using census and survey data to estimate poverty and inequality for small areas. *The Review of Economics and Statistics* **91**: 773–792.
- Østerdal LP. 2010. The mass transfer approach to multivariate discrete first order stochastic dominance: direct proof and implications. *Journal of Mathematical Economics* **46**: 1222–1228.