

Arndt, Channing; Siersbæk, Nikolaj; Østerdal, Lars Peter

Working Paper

Multidimensional first-order dominance comparisons of population wellbeing

WIDER Working Paper, No. 2015/122

Provided in Cooperation with:

United Nations University (UNU), World Institute for Development Economics Research (WIDER)

Suggested Citation: Arndt, Channing; Siersbæk, Nikolaj; Østerdal, Lars Peter (2015) : Multidimensional first-order dominance comparisons of population wellbeing, WIDER Working Paper, No. 2015/122, ISBN 978-92-9256-011-9, The United Nations University World Institute for Development Economics Research (UNU-WIDER), Helsinki, <https://doi.org/10.35188/UNU-WIDER/2015/011-9>

This Version is available at:

<https://hdl.handle.net/10419/129471>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

UNU-WIDER

30 YEARS OF RESEARCH
FOR DEVELOPMENT

WIDER Working Paper 2015/122

**Multidimensional first-order dominance
comparisons of population wellbeing**

Channing Arndt,¹ Nikolaj Siersbæk,² and Lars Peter Østerdal²

December 2015

Abstract: In this paper, we convey the concept of first-order dominance (FOD) with particular focus on applications to multidimensional population welfare comparisons. We give an account of the fundamental equivalent definitions of FOD, illustrated with simple numerical examples. An implementable method for detecting dominances is explained along with a bootstrapping procedure that yields additional information relative to what can be obtained from dominance comparisons alone. We discuss strengths and weaknesses of FOD, compared to other multidimensional population comparison concepts, and describe practical tools that enable the reader to easily use it.

Keywords: multidimensional welfare, multidimensional wellbeing, first-order dominance

JEL classification: I31, I32, D30

Figures and Tables: Provided at the end of the paper

¹UNU-WIDER, Helsinki, Finland, corresponding author: channing@wider.unu.edu; ²Department of Business and Economics and the Centre of Health Economics Research (COHERE), University of Southern Denmark.

This study has been prepared within the UNU-WIDER project '[Reconciling Africa's Growth, Poverty and Inequality Trends: Growth and Poverty Project \(GAPP\)](#)'.

Copyright © UNU-WIDER 2015

ISSN 1798-7237 ISBN 978-92-9256-011-9

Typescript prepared by Anna-Mari Vesterinen at UNU-WIDER.

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

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.

UNU-WIDER, Katajanokanlaituri 6 B, 00160 Helsinki, Finland, wider.unu.edu

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

A central question in applied welfare economics is how to make comparisons of population wellbeing across groups or over time. Appropriate comparison concepts have many potential uses. For example, if a study is able to detect that one population group is clearly worse off than another (i.e. is overall poorer or has less social welfare), society might wish to undertake policies aimed at narrowing this gap. Also, since reducing poverty or improving social welfare over time is often a key objective for public policies and reforms, the ability to make relevant comparisons over time is crucial for the formulation of meaningful goals and for policy evaluation.

The traditional approach to comparing population wellbeing is the use of a social welfare (or poverty) measure based on a one-dimensional individual wellbeing indicator, typically a monetary variable such as income or wealth. However, it has long been recognized that poverty and wellbeing are multidimensional phenomena, which are not adequately represented by a single income variable. As Sen (1976) points out, there is good reason to think that sometimes a richer person may have lower wellbeing than a poorer person, e.g. if he is disabled. This has given rise to numerous proposals of appropriate dimensions to include in multidimensional welfare analyses, including (but not restricted to) health and education (World Bank 1990) as well as standards of living (Sen 1988) to name a few.

Multidimensional welfare is often measured by aggregating multiple dimensions and weighting each dimension (see e.g. Alkire and Foster 2011; Roelen and Gassmann 2008; Rippin 2010). The weighting is primarily made in order to reflect societal judgments about different dimensions as well as to be able to obtain a single measure of the welfare for a given population. This aggregation procedure enables the analyst to rank the populations. Furthermore, the approach is very convenient and can easily be justified when there exists a reasonably high degree of consensus about which weights should be applied. There is, however, no natural and generally agreed methodology to obtaining these weights. Often, it is not easy to say if one dimension is more important than another, and, even when it is, quantifying by how much is often very difficult and perhaps not even meaningful to people.

The challenges described above have motivated the development of methods for comparing population wellbeing, poverty, or inequality with multidimensional indicators that are methodologically ‘robust’ in the sense that the conclusions obtained do not rely on predetermined weights on each dimension. In the context of applied welfare economics, such methods were popularized by Atkinson and Bourguignon (1982) who showed how stochastic dominance techniques for comparisons of probability distributions can be used to make comparisons of populations across broad classes of underlying social welfare functions. Such techniques have been further refined and applied by, e.g., Atkinson and Bourguignon (1987), Bourguignon (1989), Atkinson (1992), Bourguignon and Chakravarty (2003), Duclos et al. (2006, 2007), Gravel et al. (2009), Gravel and Mukhopadhyay (2010), Muller and Trannoy (2011), Gravel and Moyes (2012), and many others.

These contributions apply dominance concepts that rely on assumptions that are typically formulated in terms of a specified sign on the second (and higher) order partial- or cross-derivatives of the underlying individual utility function considered by a utilitarian planner. This leads to so-called lower- or upper-orthant dominance (or even more demanding concepts). For example, Duclos and Échevin (2011) assume substitutability between health and income, i.e. an underlying utility function with a negative cross partial derivative between health and income.

These concepts, while considerably more robust than methods relying on given weights, do not apply to ordinal data, where only the ranking of outcomes along each dimension is known to the analyst (based on a more-is-better assumption) but no information is available regarding for example the complementarity/substitutability relationship across the dimensions. However, data is often ordinal in nature. For example, the child is enrolled in school or not. Household sanitation facilities are adequate or not. The diet consumed by the household yesterday was adequate to meet nutritional needs or not. These are all examples of ordinal data with binary outcomes where, in each dimension, one outcome is unambiguously preferred to the other.

A natural concept for making comparisons of population distributions with multidimensional ordinal data is first-order dominance (FOD), also known as the *usual (stochastic) order* in the probability theory literature (see e.g. Shaked and Shanthikumar 2007). A finite (population, probability) distribution A first-order dominates distribution B if one can obtain distribution B from A by shifting (population, probability) mass within A from preferred to less preferred outcomes (where a less preferred outcome is not better in any dimension and is strictly worse in at least one dimension). Hence, if one distribution first-order dominates another, it is unambiguously better than the other. Thus, under the assumption that outcomes within each distribution can be ranked (e.g. we prefer the child attending school as opposed to not), the FOD approach provides a maximally robust way of making comparisons of multidimensional welfare. Technically, it does so without making any assumptions on utility functions and/or social welfare functions other than a more-is-better assumption. No additional assumptions are required about the strength of preferences for each dimension, nor about the relative desirability of changes between levels within or between dimensions (Arndt et al. 2012).

The absence of restrictive assumptions in the FOD approach makes the concept intuitively appealing. However, maximal robustness comes at a cost. First, the result of comparing two distributions may be indeterminate. In other words, it may happen that distribution A does not dominate B and B does not dominate A . This makes the analyst unable to distinguish groups A and B according to wellbeing based on the selected indicators. Second, the FOD approach provides no information about whether a dominating distribution is slightly or substantially better than a dominated distribution. This paper will discuss a way of mitigating these costs by applying a bootstrapping approach that provides a measure for the probability of observing dominances under resampling. This can serve both as a robustness check for the magnitude of the dominances observed and for the probability of observing dominance. Furthermore, if one is willing to accept the likelihood of performing well in head-to-head comparisons with other groups as an indicator of the relative wellbeing of a group, a full ranking of the groups can be calculated (Arndt et al. 2016).

The remainder of this paper is structured as follows: Section 2 provides an overview of the theory of FOD with definitions and intuitive explanations using examples. Section 3 follows up with tools for implementing the procedure in practice. Section 4 discusses faster checking algorithms and an alternative dominance criterion than FOD and lastly, Section 5 sums up and concludes.

2 Theory and examples

This section provides the basic definitions and theory of the FOD approach, illustrated with some simple examples. Furthermore, a practical linear programming method for detecting dominances is described and the bootstrapping procedure is explained.

2.1 One-dimensional FOD

Suppose first that the outcome of interest is one-dimensional. The outcome could for example be individual income (or wealth). In this case there is a natural ordering of outcomes (assuming that a higher income is better), but only this single dimension is taken into account.

2.1.1 Notations and definitions

Let X denote a finite set of real-valued outcomes. Let the distribution of wellbeing of population A be described by a probability mass function¹ f over X , i.e. $\sum f(x) = 1$ and $\sum f(x) \geq 0$ for all $x \in X$. Similarly, let the distributions of populations B and C be described by the probability mass functions g and h respectively.

As a very simple example, suppose that there are only two possible outcomes, $X = \{0,1\}$. We will always assume that higher numbers are better, so 0 is the bad outcome ('income deprived'), and 1 is the good outcome ('not income deprived'). In this situation, a population distribution is completely described by its share of individuals being income deprived. Table 1 shows distributions for three hypothetical populations.

In the one-dimensional case, f first-order dominates g if and only if any of the following (equivalent) conditions hold:²

- a) g can be obtained from f by a finite number of shifts of mass in f from one outcome to another that is *worse*.
- b) Social welfare is at least as high for f than for g for any non-decreasing additively separable social welfare function, i.e. $\sum_{x \in X} f(x)w(x) \geq \sum_{x \in X} g(x)w(x)$ for any weakly increasing real function $w(\cdot)$.
- c) $F(x) \leq G(x)$ for all $x \in X$, where $F(\cdot)$ and $G(\cdot)$ are the cumulative distribution functions (CDFs)³ corresponding to f and g .

Condition (a) provides the perhaps most intuitive definition of FOD. It provides a natural criterion for the case where one distribution is *unambiguously* better than another. Condition (b) is a robustness property in relation to social welfare comparisons and thus provides a link to welfare economics (and to expected utility theory in the case of a probability distribution). If there is FOD, social welfare will be at least as high for the dominating population *no matter the functional form of the social welfare function* as long as $w(\cdot)$ is weakly increasing. For ordinal data, this condition on the social welfare functions simply means that outcomes can be ranked from worse to better. As noted in the Introduction, no additional assumptions are required. Condition (c) turns out to be equivalent to the first two conditions and is useful for checking FOD.

¹ A probability mass function is a function that to each outcome assigns the probability of that outcome. In the context of population comparisons, it assigns the share of the population in that outcome.

² Note that FOD is conventionally defined in the weak sense, i.e. a distribution always dominates itself.

³ CDFs express the probability that the real-valued outcomes X will have a value less than or equal to x .

2.1.2 Checking one-dimensional FOD

In the one-dimensional case, FOD can be checked in a simple and effective way with direct application of Condition (c). To illustrate, consider the CDFs $F(\cdot)$, $G(\cdot)$, and $H(\cdot)$ corresponding to the three mass functions, f , g , and h , respectively, cf. Table 1. We have $F(0) = 0.35$, $G(0) = 0.50$, $H(0) = 0.40$, and of course $F(1) = G(1) = H(1) = 1$. These are shown in Figure 1 (the black line illustrates $F(\cdot)$, the gray line illustrates $G(\cdot)$, and the dotted line illustrates $H(\cdot)$).

As can be seen from the graph, $F(x) \leq H(x) \leq G(x)$ for all $x \in X$. More precisely, $F(0) < H(0) < G(0)$ whereas $F(1) = G(1) = H(1) = 1$. Hence f dominates both g and h . Since $H(x) \leq G(x)$ for all $x \in X$, h dominates g .

Condition (a) also provides an intuitive way of explaining dominances. For example, it can be seen from Table 1 that f dominates g since g can be obtained from f by shifting mass from one outcome to another that is *worse*. More precisely, shifting 15 percentage points from (1) to (0) in f yields exactly g .

2.2 Multidimensional FOD

Now suppose that the outcome is multidimensional. In the case of two dimensions, these could, for example, be income and health. In the case of three dimensions, one may wish to add education attainment, and so on.

2.2.1 Notations and definitions

Let Y be a finite set of (multidimensional) outcomes. A distribution of wellbeing of population A is described by a probability mass function f over Y , i.e. $\sum f(y) = 1$ and $f(y) \geq 0$ for all $y \in Y$. Similarly, let the distributions of populations B and C be described by the probability mass functions g and h respectively.

To illustrate, suppose that the two dimensions each have two possible outcomes, 0 or 1. Thus, $Y = \{(0,0), (0,1), (1,0), (1,1)\}$. One dimension could be *income* (dimension I) and the other dimension could be *health* (dimension II). Then the outcome (0,0) for a person means that she is deprived in both dimensions, while the outcome (1,0) means that she is not deprived in the first dimension (I) but deprived in the second dimension (II), and so on.

Suppose that three population mass functions f , g , and h are distributed as shown in Table 2. Note that distributions over dimension I (*income*) in the right most column are identical to those in Table 1 representing a situation where the same populations are considered but now an additional dimension is taken into consideration.

In the case of multidimensional outcomes, f first-order dominates g if and only if any of the following (equivalent) conditions hold:⁴

⁴The equivalence between (B) and (C) was shown by Lehmann (1955). It was also proved independently by Levhari et al. (1975). Kamae et al. (1977) observed that the equivalence between (A) and (C) is a consequence of Strassen's Theorem (Strassen 1965). See also Østerdal (2010).

- A) g can be obtained from f by a finite number of shifts of mass from one outcome to another that is *worse*.
- B) Social welfare is weakly higher for f than for g for any non-decreasing additively separable social welfare function, i.e. $\sum_{y \in Y} f(y)w(y) \geq \sum_{y \in Y} g(y)w(y)$ for any weakly increasing real function $w(\cdot)$.
- C) $\sum_{y \in Z} g(y) \geq \sum_{y \in Z} f(y)$ for any lower comprehensive set $Z \subseteq Y$.⁵

Each condition is the natural multidimensional extension of its counterpart for the one-dimensional case. Again, note that condition (A) provides an intuitive criterion for one distribution being unambiguously better than another, (B) provides a foundation in welfare economics but one that is not conveniently amenable to testing, while (C) provides a directly testable condition that may not be particularly intuitive.

2.2.2 Checking multidimensional FOD

First, due to the intuitive nature, we appeal to condition (A). It can be seen from Table 2 that f dominates g since it is possible to obtain g from f by shifting population mass in f from better to worse outcomes. More precisely, shifting 15 percentage points of probability mass from (1,1) to (0,0) in f yields exactly g , which implies that f dominates g . The distribution f is thus unambiguously preferred to the distribution g .

However, consider f and h . Neither f dominates h , nor does h dominate f . Intuitively, this is the case since f would be better if what matters most is minimization of the share of the population who are deprived in dimensions II (health) since $f(0,0) + f(1,0) = 35 < h(0,0) + h(1,0) = 40$. On the contrary, h would be better if what matters most is maximization of the share of the population not deprived in neither dimension since $h(1,1) = 50 > f(1,1) = 40$. Consequently, no dominances are detected since no assumptions are made about the relative importance of the different dimensions.

Note that the conclusion that f does not dominate h is in contrast to the one-dimensional case. This illustrates that the conclusions might change when more dimensions are added to the analysis. It is therefore important to bear in mind that an analysis applying few welfare indicators may conclude that one population dominates another whereas a multidimensional FOD analysis of the same populations with more indicators may be indeterminate. Attention should therefore be given to include important dimensions that cover overall wellbeing reasonably well.

⁵ A set $Z \subseteq Y$ is lower comprehensive if $y \in Z$, $z \in Y$, and $z \leq y$ implies $z \in Z$.

As mentioned, condition (C) provides a direct method for checking multidimensional FOD. In our example, f dominates g if and only if the following four inequalities are jointly satisfied:⁶

- i) $g(0,0) \geq f(0,0)$
- ii) $g(0,0) + g(0,1) \geq f(0,0) + f(0,1)$
- iii) $g(0,0) + g(1,0) \geq f(0,0) + f(1,0)$
- iv) $g(0,0) + g(1,0) + g(0,1) \geq f(0,0) + f(1,0) + f(0,1)$.

Considering the distributions in Table 2 in relation to the four inequalities above, it can be seen that, when comparing f and g , each of the four inequalities (i)-(iv) are (strictly) satisfied; (i) $g(0,0) \geq f(0,0)$ since $0.25 > 0.10$, (ii) $g(0,0) + g(0,1) \geq f(0,0) + f(0,1)$ since $0.25 + 0.25 > 0.10 + 0.25$, (iii) $g(0,0) + g(1,0) \geq f(0,0) + f(1,0)$ since $0.25 + 0.25 > 0.10 + 0.25$, and (iv) $g(0,0) + g(1,0) + g(0,1) \geq f(0,0) + f(1,0) + f(0,1)$ since $0.25 + 0.25 + 0.25 > 0.10 + 0.25 + 0.25$. Hence, f dominates g as already observed. When comparing each of the other distribution pairs, it can easily be seen that at least one of the inequalities is violated. For example, when comparing f and h , inequalities (i)-(iii) are satisfied whereas inequality (iv) is not.

2.2.3 Detecting FOD in practice

Criterion (C) provides a simple method for detecting dominance, which can be visually perceived in cases with few outcomes such as in the example with two binary indicators giving four different outcomes in total. However, the number of inequalities to be checked increases drastically when more dimensions and levels are added. For real world applications, computationally efficient algorithms for checking dominance are required (Range and Østerdal 2013). Mosler and Scarsini (1991) and Dyckerhoff and Mosler (1997) show that, appealing to definition (A), checking FOD corresponds to determining if a certain linear programme has a feasible solution. The first empirical implementation of this approach was provided by Arndt et al. (2012) in a study of child poverty in Mozambique and Vietnam. FOD can thus be determined using a linear programming package.

Let \mathcal{A} and \mathcal{B} be two populations characterized by probability mass functions f and g respectively. For outcomes \mathbf{y} and \mathbf{y}' with $\mathbf{y}' \leq \mathbf{y}$, let $t_{\mathbf{y},\mathbf{y}'}$ be the amount of probability mass transferred from outcome \mathbf{y} to \mathbf{y}' . Note that the first subscript denotes the source of the transfer whereas the second denotes the destination.

Given the conditions outlined above, population \mathcal{A} dominates population \mathcal{B} if and only if there exist a feasible solution to the following linear programme:⁷

⁶ Note that the fifth inequality $g(0,0) + g(1,0) + g(0,1) + g(1,1) \geq f(0,0) + f(1,0) + f(0,1) + f(1,1)$ is always satisfied with equality by the definition of the probability mass functions since $\sum f(\mathbf{y}) = \sum g(\mathbf{y}) = 1$.

⁷ The approach differs slightly from that outlined by Mosler and Scarsini (1991) and Dyckerhoff and Mosler (1997). In particular, the transfers here are absolute and not relative. Note furthermore that most linear programming packages require the specification of an objective function. This can be defined as an arbitrary constant function.

$$f(y) + \sum_{y' \geq y} t_{y',y} - \sum_{y' \leq y} t_{y,y'} = g(y) \quad \forall y \in Y, t_{y,y'} \geq 0, t_{y,y} = 0. \quad (1)$$

To provide an example, let us return to Table 2. There, f dominates g since it is possible to obtain g from f by shifting 15 percentage points of probability mass from outcome (1,1) to (0,0), c.f. condition (A). In terms of the linear programme in (1), this implies that for $y = (0,0)$ and $y' = (1,1)$, $t_{y',y} = 0.15$ in order for the equality to be fulfilled for $y = (0,0)$. Furthermore, for $y = (1,1)$ and $y' = (0,0)$, $t_{y,y'} = 0.15$ in order for the equality to be fulfilled for $y = (1,1)$. It is thus possible to fulfill all the constraints for all $y \in Y$ whereby f dominates g .

An implementation of the FOD approach using linear programming in GAMS is reviewed in Section 3.

2.3 Mitigating the limitations of FOD

As mentioned previously, due to the absence of strong assumptions such as predetermined weights, there are some inherent limitations to the FOD approach. First, when comparing two groups, no dominances may be found and hence the FOD approach might yield an indeterminate result. This provides little information about the populations' relative wellbeing, as was the case when considering f and b with multiple dimensions in Section 2.2.2. Second, the FOD approach provides no information about the strength of dominance. For example, if population A dominates population B , the FOD check itself provides no information as to whether A is marginally or substantially better than B . Both of these limitations can be mitigated using a bootstrapping approach as described in the following sub-section.

2.3.1 Bootstrapping

To mitigate the first and second limitations mentioned above, a bootstrapping procedure can be applied (Arndt et al. 2012). In general, bootstrapping is a procedure that relies on random sampling with replacement from the original data set. When comparing population A and B , J samples of size K are drawn with replacement for each population group where $K \leq N$, N being the number of individuals in that population in the original sample. In the bootstrap procedure outlined in Section 3, the samples are drawn in clusters from each stratum with $K = N$. When a cluster is drawn, all households in that cluster are drawn. Due to the drawing with replacement, each cluster (and thus household) may appear more than once. The FOD approach is then applied to each of the J bootstrap samples. When these repeated bootstrap samples are compared using the FOD approach, the final output can be interpreted as an empirical probability that population A dominates population B since the original sample is a sub-sample from a larger population. These probabilities yield significantly more information than only applying FOD on the original data where, for example, an indeterminate result will make it impossible to draw further conclusions about the comparative wellbeing of the two populations. The bootstrapping procedure thus enables the analyst to extract some information about the strength of conclusion based on the probability of dominance under resampling.

For example, with bootstrapping, we may find that occasionally A dominates B and occasionally the inverse is true, but most of the time the results are indeterminate, i.e. *rough equality* of A and B . Alternatively, we may find that the probability that A dominates B is fairly high, the probability that B dominates A is very low or zero, and the probability of an indeterminate result is somewhat low, i.e. *likely dominance* of A over B , or we might find that A dominates B almost always, i.e. *solid dominance* of A over B .

As a concrete example, say that population A dominates B in 995 of the $J = 1,000$ bootstrap samples and that A dominates C in 870 of the bootstrap samples. This corresponds to a 99.5 per cent chance of A dominating B and a 87 per cent change of A dominating C . In this example, A is thus better than C (*likely dominance* of A over C) and considerably better than B (*solid dominance* of A over B). If no dominations between B and C are obtained in the original sample, the ranking of these is ambiguous. However, if for example B dominates C in 4 (0.4 per cent) of the bootstrap samples whereas C dominates B in 80 (8 per cent) of the bootstrap samples, C is seemingly better than B (though *rough equality* of B and C).

Furthermore, if one is willing to accept the tendency to outperform other groups as overall relative indicator of population wellbeing, it is possible to provide an intuitive ranking of all population groups via the Copeland (1951) method, which is analogous to the way in which teams are ranked by assigning points to wins, draws, and losses from matchups in a sports tournament. For instance, for each population group (n population groups in total), one can count how many of the $(n - 1)$ other population groups it dominates and from that subtract the number of times it is dominated by these other groups. This yields a score in the interval $[-(n - 1), n - 1]$ which can then be normalized to the interval $[-1, 1]$ (see e.g. Arndt et al. 2016).

3 Implementation

This section briefly describes how to implement the FOD approach for analysis across space and/or time. The GAMS code for detecting dominances and the bootstrapping procedure is available from the authors upon request. First, one requires data. This can come in the form of Demographic and Health Surveys, census data, or data from living standards measurements (LSMS) type surveys, among other possibilities.

Second, one must organize the data into populations and then into groups whose welfare levels one would like to compare. There are enormous possibilities for populations and division into groups. An example of a population to study might be children aged 0-5. This population could be grouped by gender, province, and time period. One would then be setting up to examine whether, for example, girls aged 0-5 in province A at time T are better off than boys aged 0-5 in province A at time T alongside many other possible permutations. A second population example could be households and subsequent groups could be determined by the ethnicity of the household head. One would then be setting up to examine whether households headed by ethnic group A are better off than households headed by ethnic group B , and so forth.

Third, one must define welfare indicators. As noted, proper definition of indicators is critical. The indicators must apply to the population in question. If the population is children aged 0-5, then an indicator like school attendance is not relevant because children that young typically do not attend school. School attendance would more properly apply to the population of children aged 7-17. For children aged 0-5, relevant indicators often include anthropometric data, education level of the mother or primary caretaker, vaccination record, proximity and/or use of health services, and other similar indicators.

It is important to highlight that, while the FOD procedure admits in principle ordinal data, in practice, the procedure in GAMS is coded to consider only binary indicators. Hence, one might classify children aged 0-5 who are any one of stunted, under-weight, or wasted as deprived (0) in the anthropometrics dimension and all others as not deprived (1). Similarly, one might consider children whose primary caretaker has completed at least primary school as not deprived and all others as deprived. These cutoff points (e.g., complete primary school) are often important

determinants of results and should be considered carefully. The end result should be a data set where each member of the population is an observation. Variables identify the group to which this population member belongs and the welfare status (deprived or not deprived) of the population member for each chosen welfare indicator.

Fourth, one must consider the number of indicators to employ for the analysis. With binary indicators, the number of permutations of welfare states increases by a factor of two with the addition of each indicator. Specifically, there are 2^N permutations where N is the number of indicators chosen. There are trade-offs here. More indicators imply a broader analysis. At the same time, more indicators lead to many more permutations eventually resulting in a very small number of observations occupying certain permutations even for very large data sets. For example, if one chooses to make comparisons on the basis of seven indicators, there are then $2^7 = 128$ possible permutations. Example permutations include those deprived in all indicators (0,0,0,0,0,0,0), those not deprived in all indicators (1,1,1,1,1,1,1), those not deprived in indicator one and deprived in all other indicators (1,0,0,0,0,0,0) and so forth (125 additional permutations). Because of this ‘curse of dimensionality’, the available code handles only up to seven indicators. Further, because we are often making comparisons between specific sub-groups (e.g., children aged 0-5 in region A at time T versus children aged 0-5 in region B at time T), the number of observations for these groups may be insufficient to adequately populate 128 permutations. In practice, five dimensions are often chosen resulting in $2^5 = 32$ permutations.⁸

Once the number of binary indicators has been chosen, it is straightforward to collapse the data to show the shares of each group by permutation. So, if a group is girls aged 0-5 in region A at time T and the number of welfare indicators is 5, then the sum of shares across the 32 welfare permutations should be equal to one. This procedure generates the distributions for each group.

Finally, these distributions are then fed into the GAMS programme that compares groups to one another. This generates one set of comparisons using the original data. Codes are also in place to draw bootstrap samples from the original data in order to run multiple comparisons and thus estimate a probability of domination. The bootstrap procedure is applied to the data generated at the third step discussed above. The GAMS programme uses the linear programming approach outlined in Section 2.2.3 to detect dominance. A series of outputs are also automatically generated including scores and rankings based on a Copeland approach.

4 Further considerations

4.1 Faster solution algorithms

The linear programming approach presented provides a practical method for checking FOD for many applied problems. For most applications, the method is computationally fast enough to allow for a great number of pairwise comparisons as well as, if desired, hundreds of bootstrap

⁸ If the dimensions can be naturally grouped, one may initially collapse related dimensions into a single dimension and then (sequentially) refine dimensions. For more details and an illustration of such an approach, see Hussain et al. (2015).

repetitions for each pair of distributions compared. In applications with multiple binary indicators, the linear programming method is particularly suitable.⁹

However, the linear programming approach is not the fastest possible way to make FOD comparisons. Using a network flow formulation of the problem, as outlined in Preston (1974) or Hansel and Troallic (1978), it is possible to check FOD via computation of the maximum flow. As discussed in Range and Østerdal (2013), the problem of checking FOD for multidimensional distributions can also be formulated as a special bipartite network problem related to the classical transportation problem. Generally, these formulations are computationally faster than the linear programming method. In particular, for the bivariate case, Range and Østerdal (2013) provide an algorithm for checking FOD where the worst-case computational complexity grows linearly in the size of the problem (determined by the total number of outcomes).

4.2 Alternative dominance criterion

There are many other dominance criteria in the literature than FOD. In general, the alternative dominance criteria all impose stronger underlying assumptions on the underlying utility/social welfare functions. A comprehensive overview of alternative dominance criteria is outside the scope of this paper (see e.g. Shaked and Shanthikumar 2007, for an extensive review). However, we will compare FOD with the lower-orthant dominance ordering, which is one of the most frequently used alternative dominance criteria in welfare economics.

4.2.1 Orthant stochastic orderings

As mentioned in the Introduction, the FOD approach differs from the criteria for robust welfare comparisons of the Atkinson-Bourguignon type (see Atkinson and Bourguignon 1982; Atkinson and Bourguignon 1987; Bourguignon 1989; Atkinson 1992). These are variations of *orthant stochastic orderings* (see Dyckerhoff and Mosler 1997) even though the name first-order dominance has sometimes been used synonymously with orthant stochastic orderings in the welfare economics literature (e.g. Atkinson and Bourguignon 1982). If one assumes substitutability between dimensions (as e.g. Duclos and Échevin 2011, where substitutability between health and income is assumed, i.e. an underlying utility function with a negative cross partial derivative), a criterion less restrictive than FOD can be used. In particular, *f* orthant dominates *g* if and only if

$$C_0) \quad \sum_{z \leq y} g(z) \geq \sum_{z \leq y} f(z) \text{ for all } y \in Y.$$

The label (C_0) is used to indicate that this condition relates to condition (C) in the case of multidimensional FOD in Section 2.2.1. However, condition (C_0) is less restrictive than (C). This implies that condition (C_0) may be satisfied even though conditions (A), (B), and (C) are not. For a two-dimensional comparison with binary indicators (as in Table 2), *f* orthant dominates *g* if and only if each of the following three conditions are satisfied:

- $i_0) \quad g(0,0) \geq f(0,0)$
- $ii_0) \quad g(0,0) + g(0,1) \geq f(0,0) + f(0,1)$
- $iii_0) \quad g(0,0) + g(1,0) \geq f(0,0) + f(1,0).$

⁹ Even with binary indicators, the linear programming approach might actually be computationally challenged, but that is only so if the number of dimensions is large (Hussain et al. 2015).

The labels i_0 - iii_0 are used to indicate that these conditions relate to conditions i - iii in Section 2.2.2. Note that the fourth inequality (iv) need not be satisfied for orthant dominance. Returning to f and b in Table 2, recall that neither f first-order dominates b , nor does b first-order dominate f . However, in the case of orthant dominance, it can be seen that; (i_0) $h(0,0) \geq f(0,0)$ since $0.30 > 0.10$, (ii_0) $h(0,0) + h(0,1) \geq f(0,0) + f(0,1)$ since $0.30 + 0.10 > 0.10 + 0.25$, and (iii_0) $h(0,0) + h(1,0) \geq f(0,0) + f(1,0)$ since $0.30 + 0.10 > 0.10 + 0.25$. Hence f orthant dominates b even though f does *not* first-order dominate b .

5 Conclusion

Population wellbeing is increasingly recognized as a multidimensional phenomenon that is not adequately described by a single dimension (e.g. by income only). Several methods of measuring and comparing welfare have been proposed where application of a weighting- or counting scheme to different dimensions is used. It is often, however, difficult to determine these weights. Due to the sensitivity of the outcome to the weights applied, different conclusions about welfare rankings are likely to occur if the weighting scheme differs from one analysis to another. While comparisons using, for example, lower orthant (stochastic) orderings following Atkinson and Bourguignon (1982) are considerably more ‘robust’ than applying weighting schemes, they typically apply conditions formulated in terms of the second (or higher) order partial cross derivative and do not apply to ordinal data.

The first-order (stochastic) dominance (FOD) approach requires only that the outcomes in each dimension can be ranked from worse to better. The FOD approach can be applied to ordinal multidimensional data, enabling the analyst to perform wellbeing comparisons across population groups with a minimum of assumptions imposed. FOD is thus robust across all possible weighting schemes. This advantage is accompanied by limitations in that the FOD approach can yield indeterminate outcomes and does not directly provide information with respect to degree of dominance. A bootstrapping approach can be used to obtain more information thus mitigating these limitations. Moreover, a Copeland approach can be used to obtain a ranking (i.e. a complete and transitive ordering) of all groups being compared.

Finally, it is worth mentioning that even though the FOD approach and bootstrapping procedure enable the analyst to rank population welfare without assumptions about weights, the analysis should ideally be performed together with alternative welfare measurements, which provide cardinal information about the relative size of wellbeing differences under fixed weights assumptions. As Ferreira (2011) puts it, looking at a few core, truly irreducible, dimensions and applying dominance analysis (as well as a number of indicies), is likely to contribute to the design and targeting of policy actions. Thus, FOD comparisons should form a part of a broader population wellbeing analysis strategy.

References

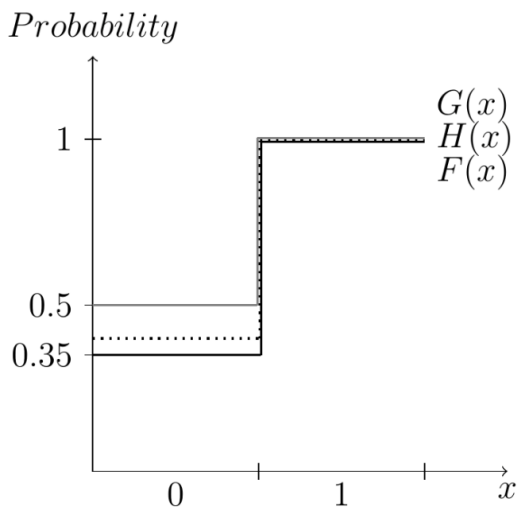
- Alkire, S., and J. Foster (2011). ‘Counting and multidimensional poverty measurement’. *Journal of Public Economics*, 95(7-8): 476-87.
- Arndt, C., R. Distanto, M.A. Hussain, L.P. Østerdal, P.L. Huong, and M. Ibraimo (2012). ‘Ordinal welfare comparisons with multiple discrete indicators: A first order dominance approach and application to child poverty’. *World Development*, 40(11): 2290-301.

- Arndt, C., M.A. Hussain, S. Vincenzo, F. Tarp, and L.P. Østerdal (2016). 'Poverty mapping based on first order dominance with an example from Mozambique'. *Journal of International Development*, (forthcoming). DOI: 10.1002/jid.3200
- Atkinson, A.B. (1992). 'Measuring poverty and differences in family composition'. *Economica*, 59(233): 1-16.
- Atkinson, A.B., and F. Bourguignon (1982). 'The comparison of multi-dimensioned distributions of economic status'. *Review of Economic Studies*, 42(2): 183-201.
- Atkinson, A.B., and F. Bourguignon (1987). 'Income distribution and differences in need'. In G.R. Feiwel (ed.), *Arrow and the foundations of the theory of economic policy*. New York: New York University Press.
- Bourguignon, F. (1989). 'Family size and social utility: Income distribution dominance criteria'. *Journal of Econometrics*, 42(1): 67-80.
- Bourguignon, F., and S.R. Chakravarty (2003). 'The measurement of multidimensional poverty'. *Journal of Economic Inequality*, 1(1): 25-49.
- Copeland, A.H. (1951). *A 'Reasonable' Social Welfare Function*. University of Michigan Seminar on Applications of Mathematics to the Social Sciences.
- Duclos, J.-Y., D.E. Sahn, and S.D. Younger (2006). 'Robust multidimensional poverty comparisons'. *Economic Journal*, 116(514): 943-68.
- Duclos, J.-Y., D.E. Sahn, and S.D. Younger (2007). 'Robust multidimensional poverty comparisons with discrete indicators of well-being'. In S.P. Jenkins and J. Micklewright (eds), *Inequality and Poverty Re-examined*. Oxford: Oxford University Press.
- Duclos, J.-Y., and D. Échevin (2011). 'Health and income: A robust comparison of Canada and the US'. *Journal of Health Economics*, 30(2): 293-302.
- Dyckerhoff, R., and K. Mosler (1997). 'Orthant orderings of discrete random vectors'. *Journal of Statistical Planning and Inference*, 62(2): 193-205.
- Ferreira, F.H.G. (2011). 'Poverty is multidimensional. But what are we going to do about it?'. *Journal of Economic Inequality*, 9(3): 493-95.
- Gravel, N., and P. Moyes (2012). 'Ethically robust comparisons of bidimensional distributions with an ordinal attribute'. *Journal of Economic Theory*, 147(4): 1384-426.
- Gravel, N., P. Moyes, and B. Tarroux (2009). 'Robust international comparisons of distributions of disposable income and regional public goods'. *Economica*, 76(303): 432-61.
- Gravel, N., and A. Mukhopadhyay (2010). 'Is India better off today than 15 years ago? A robust multidimensional answer'. *Journal of Economic Inequality*, 8(2): 173-95.
- Hansel, G., and J. Trollic (1978). 'Mesures marginales et théorème de Ford-Fulkerson'. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 43(3): 245-51.
- Hussain, M.A., M.M. Jørgensen, and L.P. Østerdal (2015). 'Refining population health comparisons: A multidimensional first order dominance approach'. *Social Indicators Research*, (forthcoming). DOI: 10.1007/s11205-015-1115-2
- Kamae, T., U. Krengel, and G.L. O'Brien (1977). 'Stochastic inequalities on partially ordered spaces'. *Annals of Probability*, 5(6): 899-912.
- Lehmann, E.L. (1955). 'Ordered families of distributions'. *Annals of Mathematical Statistics*, 26(3): 399-419.

- Levhari, D., J. Paroush, and B. Peleg (1975). 'Efficiency analysis for multivariate distributions'. *Review of Economic Studies*, 42(1): 87-91.
- Mosler, K.C., and M.R. Scarsini (1991). *Stochastic orders and decision under risk*. Heyward, CA: Institute of Mathematical Statistics.
- Muller, C., and A. Trannoy (2011). 'A dominance approach to the appraisal of the distribution of well-being across countries'. *Journal of Public Economics*, 95(3-4): 239-46.
- Preston, C. (1974). 'A generalization of the FKG inequalities'. *Communications in Mathematical Physics*, 36(3): 233-41.
- Range, T.M., and L.P. Østerdal (2013). 'Checking bivariate first order dominance'. Discussion Papers on Business and Economics 9/2013. Odense: Department of Business and Economics, University of Southern Denmark.
- Rippin, N. (2010). 'Poverty severity in a multidimensional framework: The issue of inequality between dimensions'. Courant Research Centre, Poverty, Equity and Growth Discussion Paper 47. Göttingen: University of Göttingen.
- Roelen, K., and F. Gassmann (2008). 'Measuring child poverty and wellbeing: A literature review'. Maastricht Graduate School of Governance Working Paper Series 2008/WP001. Maastricht: Maastricht University.
- Sen, A.K. (1976). 'Poverty: An ordinal approach to measurement'. *Econometrica*, 44(2): 219-31.
- Sen, A.K. (1988). *The Standard of Living*. Cambridge: Cambridge University Press.
- Shaked, M., and J.G. Shanthikumar (2007). *Stochastic orders*. New York: Springer Science and Business Media.
- Strassen, V. (1965). 'The existence of probability measures with given marginals'. *Annals of Mathematical Statistics*, 36(2): 423-39.
- World Bank (1990). *World Development Report 1990: Poverty*. Oxford University Press for the World Bank.
- Østerdal, L.P. (2010). 'The mass transfer approach to multivariate discrete first order stochastic dominance: Direct proof and implications'. *Journal of Mathematical Economics*, 46(6): 1222-28.

Figures

Figure 1: CDFs, $F(x)$, $G(x)$, and $H(x)$, one-dimensional



Source: Authors' hypothetical example from Table 1.

Tables

Table 1: Distributions f , g , and h (per cent), one-dimensional

f		Total
Income	0 (deprived)	35
	1 (not deprived)	65
Total		100
Population B		
g		Total
Income	0 (deprived)	50
	1 (not deprived)	50
Total		100
Population C		
h		Total
Income	0 (deprived)	40
	1 (not deprived)	60
Total		100

Source: Authors' hypothetical example.

Table 2: Distributions f , g , and h (per cent), two-dimensional

f		Population A		Total
		II (health)		
		0 (deprived: bad)	1 (not deprived: good)	
I (income)	0 (deprived: poor)	10	25	35
	1 (not deprived: rich)	25	40	65
Total		35	65	100

g		Population B		Total
		II (health)		
		0 (deprived: bad)	1 (not deprived: good)	
I (income)	0 (deprived: poor)	25	25	50
	1 (not deprived: rich)	25	25	50
Total		50	50	100

h		Population C		Total
		II (health)		
		0 (deprived: bad)	1 (not deprived: good)	
I (income)	0 (deprived: poor)	30	10	40
	1 (not deprived: rich)	10	50	60
Total		40	60	100

Source: Authors' hypothetical example.