

Mehdi, Tahsin

Article

Testing for stochastic dominance up to a common relative poverty line

Econometrics

Provided in Cooperation with:

MDPI – Multidisciplinary Digital Publishing Institute, Basel

Suggested Citation: Mehdi, Tahsin (2020) : Testing for stochastic dominance up to a common relative poverty line, *Econometrics*, ISSN 2225-1146, MDPI, Basel, Vol. 8, Iss. 1, pp. 1-5, <https://doi.org/10.3390/econometrics8010005>

This Version is available at:

<https://hdl.handle.net/10419/247553>

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.



<https://creativecommons.org/licenses/by/4.0/>

Article

Testing for Stochastic Dominance up to a Common Relative Poverty Line

Tahsin Mehdi 

Department of Economics, Ryerson University, Toronto, ON M5B 2K3, Canada; tahsin.mehdi@ryerson.ca

Received: 18 October 2019; Accepted: 5 February 2020; Published: 11 February 2020



Abstract: Although a wide array of stochastic dominance tests exist for poverty measurement and identification, they assume the income distributions have independent poverty lines or a common absolute (fixed) poverty line. We propose a stochastic dominance test for comparing income distributions up to a common relative poverty line (i.e., some fraction of the pooled median). A Monte Carlo study demonstrates its superior performance over existing methods in terms of power. The test is then applied to some Canadian household survey data for illustration.

Keywords: stochastic dominance; common poverty line; relative poverty line; pooled quantile; bootstrap inference

JEL Classification: C01; C12; I32

1. Introduction

The seminal works of (Sen 1976; Foster et al. 1984; Atkinson 1987) have propagated a growing body of literature surrounding poverty measurement. While earlier works emphasize the identification of poverty, later works highlight the need for accurate and reliable statistical inference (e.g., (Kakwani 1993; Davidson and Duclos 2000; Zheng 2001; Thompson 2012; Mehdi 2017)). From a research as well as a policy standpoint, there is always interest in comparing poverty outcomes between different income distributions. However, as pointed out by (Garcia-Gomez et al. 2019), analysts often face the multiplicity of poverty indices problem. Since poverty measures depend on a poverty line, which is an income threshold dividing the poor and non-poor, distributional orderings are sensitive to the choice of poverty lines (i.e., rank reversals could occur when switching poverty lines).

To overcome the multiplicity of poverty indices problem, researchers appeal to stochastic dominance which is a robust distributional ranking method with wide-ranging applications in multiple fields. Rather than relying on estimates from singular points in distributions, stochastic dominance examines the entire support by considering the cumulative distribution function (CDF). In the context of poverty comparisons, stochastic dominance permits partial orderings of income distributions by considering all permissible income thresholds within a pre-specified range of poverty lines. If a dominance relation can be established between distributions, debates surrounding the choice of poverty lines can effectively be avoided since one distribution would always exhibit lower poverty regardless of the poverty line (see (Atkinson 1987)).

Several statistical tests of stochastic dominance have been put forth in the literature (e.g., (McFadden 1989; Kaur et al. 1994; Anderson 1996; Davidson and Duclos 2000; Barrett and Donald 2003; Barrett et al. 2016; Thompson and Stengos 2012)), but the ones geared towards poverty measurement assume either separate poverty lines for each distribution or a common absolute poverty line between distributions. As (Thompson 2012) points out, if there is a situation where we may be interested in comparing poverty outcomes between subgroups of the same population

(e.g., males and females), setting two separate poverty lines may lead to rather incongruous findings.¹ With increasing usage of relative poverty measures by international organizations (e.g., (OECD 2016)), we develop the asymptotic framework for testing for stochastic dominance up to a common relative poverty line (i.e., some fraction of the pooled median income level of two distributions).

(Barrett and Donald 2003) propose a test for stochastic dominance based on the one-sided Kolmogorov–Smirnov type test statistic which considers the supremum of the distances between the CDFs. Our test, on the other hand, is similar to that of (Davidson and Duclos 2000) in the sense that it relies on evaluating a finite number of distances between the distributions throughout the support. Our proposed test, much like those of (Anderson 1996; Davidson and Duclos 2000; Thompson and Stengos 2012), suffers from the issue of inconsistency due to the fact that such tests rely on examining the CDFs at a finite number of points as opposed to those put forth by (Barrett and Donald 2003; McFadden 1989; or Kaur et al. 1994) which consider either the supremum or infimum of the distances (see, e.g., (Davidson and Duclos 2000; Thompson and Stengos 2012)). However, the advantage offered by the former type of tests is that they make use of the covariances between estimates at the different points of the CDFs, which in theory leads to increased statistical power.

The remainder of this article is organized as follows. Section 2 briefly discusses the notion of stochastic dominance and its relation to poverty measures. Section 3 derives the asymptotic framework of our proposed test for stochastic dominance when the poverty line is some fraction of the pooled median income level. Section 4 presents a Monte Carlo study to assess the size and power of the test. Section 5 illustrates the proposed test using Canadian household survey data. Section 6 provides the conclusions.

2. Stochastic Dominance and Poverty Measurement

Consider two income distributions (or some other measure of individual welfare), characterized by CDFs, F_A and F_B , with support contained in the non-negative real number line. Following similar notation as (Davidson and Duclos 2000), let

$$D_A^s(x) = \frac{1}{(s-1)!} \int (x-y)_+^{s-1} dF_A(y), \quad (1)$$

where $s \geq 1$, and $(x-y)_+^{s-1} = (x-y)^{s-1} I(y \leq x)$, where $I(\cdot)$ is an indicator function that equals 1 if its argument is true, and 0 otherwise. It is straightforward to check that $D_A^1(x) = F_A(x)$. If a poverty line z is established, an individual with income y is said to be poor if $y \leq z$ (this follows from the so-called “focus axiom”; see, e.g., (Foster 1984)). Thus, $F_A(z)$ measures the proportion of individuals in subgroup A below the poverty line (also known as the headcount ratio). Let $D_B^s(x)$ be defined analogously. The D_A^1 curve is typically referred to as the poverty incidence curve, D_A^2 is the poverty deficit curve, and D_A^3 is the poverty severity curve (see, e.g., (Ravallion 1994)). Distribution A is said to stochastically dominate B at order s up to poverty line z if $D_A^s(x) \leq D_B^s(x) \forall x \leq z$. First-order stochastic dominance (i.e., $s = 1$) guarantees dominance at higher orders (i.e., if $D_A^1(x) \leq D_B^1(x) \forall x \leq z$, then $D_A^s(x) \leq D_B^s(x) \forall x \leq z, s > 1$; see, e.g., Lemma 1, (Davidson and Duclos 2000)).

The notion of stochastic dominance has broader implications for popular classes of poverty indices such as those proposed by (Foster et al. 1984): $P_\gamma = \int [(z-y)/z]^\gamma I(y \leq z) dF(y)$, where $\gamma \geq 0$ is a poverty “aversion” parameter. Thus, the class of indices is based on the normalized poverty gap, $(z-y)/z$, or income shortfall as a share of the poverty line, of the poor. It is easy to see that, when $\gamma = 0$, the index simply becomes the headcount ratio (i.e., $P_0 = F(z)$), which measures the proportion of the population below the poverty line. Thus, if indeed $D_A^1(x) \leq D_B^1(x) \forall x \leq z$, then the implication

¹ (Thompson 2012) proposes an empirical likelihood-based test for comparing poverty measures between two distributions using a poverty line set to some fraction of the pooled median of the combined distribution, but the method cannot detect stochastic dominance since it permits only equality restrictions on the hypotheses.

is that P_γ will not only show lower poverty incidence for distribution A for all poverty lines up to z , but lower poverty deficit, lower poverty severity, etc.

Consider a population of size $N = N_A + N_B$ composed of N_A individuals from subgroup A , and N_B individuals from subgroup B . Let $F(x) = w_A F_A(x) + w_B F_B(x)$ be the pooled income distribution where $w_A = N_A/N$ and $w_B = N_B/N$. Assume a pooled quantile-based poverty line set at $z = c\zeta_q$ for some fraction $c \in [0, 1]$, where ζ_q is a quantile of order q (i.e., $F(\zeta_q) = q$).

Checking for restricted stochastic dominance is tantamount to examining the differences between the distributions at all points leading up to z . We follow (Davidson and Duclos 2000) and construct test statistics at equidistant grid points that lie below z . Consider J grid points $\{\alpha_i \zeta_q\}_{i=1}^J$ where $\alpha_i \in [0, c]$.² Let the vector of differences between the two distributions at the J grid points be given by $\Delta^s = (\Delta_1^s, \dots, \Delta_J^s) = (D_{B,1}^s - D_{A,1}^s, \dots, D_{B,J}^s - D_{A,J}^s)$.

3. Estimation and Inference

Let $\{y_{A,i}\}_{i=1}^{n_A}$ and $\{y_{B,i}\}_{i=1}^{n_B}$ be random iid draws from F_A and F_B , respectively, and let $n = n_A + n_B$ be the pooled sample size. Assume that $n \rightarrow \infty$ implies $n_A \rightarrow \infty$ and $n_B \rightarrow \infty$, and n_A/N_A and n_B/N_B are sufficiently small so that no finite population adjustment is necessary. At the j th grid point, D_A^s can be consistently estimated by

$$\hat{D}_{A,j}^s = \frac{1}{n_A(s-1)!} \sum_{i=1}^{n_A} (\hat{z}_j - y_{A,i})_+^{s-1}, \tag{2}$$

where $\hat{z}_j = \alpha_j y_{(r)}$, $y_{(r)}$ is the r th order statistic of the pooled sample $y = (y^A, y^B)'$ with $r = [nq]$, and $[nq]$ is the integer part of nq . For subgroup B , D_B^s can be estimated in a similar manner.

If F_A and F_B are differentiable and have finite first two moments, then Δ^s can be consistently estimated by $\hat{\Delta}^s = (\hat{\Delta}_1^s, \dots, \hat{\Delta}_J^s) = (\hat{D}_{B,1}^s - \hat{D}_{A,1}^s, \dots, \hat{D}_{B,J}^s - \hat{D}_{A,J}^s)$. Using similar arguments as ((Zheng 2001), Section 4.2), an asymptotic expression for the j th difference is given by

$$\Delta_j^s = \frac{1}{n_B(s-1)!} \sum_{i=1}^{n_B} (z_j - y_{B,i})_+^{s-1} - \frac{1}{n_A(s-1)!} \sum_{i=1}^{n_A} (z_j - y_{A,i})_+^{s-1} + \alpha_j g_j(y_{(r)} - \zeta_q) + o_p(n^{-1/2}), \tag{3}$$

where $g_j = a_j^B - a_j^A + I(s=1)[f_B(z_j) - f_A(z_j)]$, $a_j^A = \partial D_{A,j}^s / \partial z_j$, and f_A is the underlying density function of distribution A .

Using the Bahadur representation (see, e.g., (Zheng 2001, p. 351)), we can express the difference between the sample quantile and population quantile as

$$y_{(r)} - \zeta_q = \frac{q - \hat{F}(\zeta_q)}{f(\zeta_q)} + o_p(n^{-1/2}), \tag{4}$$

where $\hat{F}(\zeta_q) = n^{-1} \sum_{i=1}^n I(y_i \leq \zeta_q)$, and $f(\zeta_q) = w_A f_A(\zeta_q) + w_B f_B(\zeta_q)$ is the underlying population density function. Thus, (3) becomes

$$\Delta_j^s = \frac{1}{n_B(s-1)!} \sum_{i=1}^{n_B} (z_j - y_{B,i})_+^{s-1} - \frac{1}{n_A(s-1)!} \sum_{i=1}^{n_A} (z_j - y_{A,i})_+^{s-1} - \frac{\alpha_j g_j}{f(\zeta_q)} \hat{F}(\zeta_q) + \frac{\alpha_j g_j q}{f(\zeta_q)} + o_p(n^{-1/2}). \tag{5}$$

² For instance, if the poverty line is set to 50% ($c = 0.5$) of the pooled median ($q = 0.5$), then some possible grid points could be 10%, 20%, 30%, 40%, and 50% of the pooled median.

Let the joint population moments of order $2s - 2$ of y_A and y_B be finite and suppose that F_A and F_B are differentiable. Then, $\sqrt{n}(\hat{\Delta}^s - \Delta^s)$ will converge in distribution to a normal random vector with mean vector zero and covariance matrix Σ with typical element

$$\begin{aligned} \text{Cov}(\hat{\Delta}_j^s, \hat{\Delta}_k^s) = & \left[\frac{1}{[(s-1)!]^2} E_A[(z_j - y)_+^{s-1} (z_k - y)_+^{s-1}] - D_{A,j}^s D_{A,k}^s \right] / w_A \\ & + \left[\frac{1}{[(s-1)!]^2} E_B[(z_j - y)_+^{s-1} (z_k - y)_+^{s-1}] - D_{B,j}^s D_{B,k}^s \right] / w_B \\ & - \alpha_k g_k [D_{B,j}^s [1 - F_B(\xi_q)] - D_{A,j}^s [1 - F_A(\xi_q)]] / f(\xi_q) \\ & - \alpha_j g_j [D_{B,k}^s [1 - F_B(\xi_q)] - D_{A,k}^s [1 - F_A(\xi_q)]] / f(\xi_q) \\ & + \alpha_j \alpha_k g_j g_k [w_A F_A(\xi_q) [1 - F_A(\xi_q)] \\ & + w_B F_B(\xi_q) [1 - F_B(\xi_q)]] / f^2(\xi_q), \forall j, k. \end{aligned}$$

In practice, Σ can be consistently estimated by $\hat{\Sigma}$ which will have typical element

$$\begin{aligned} \hat{\text{Cov}}(\hat{\Delta}_j^s, \hat{\Delta}_k^s) = & \left[\frac{1}{n_A [(s-1)!]^2} \sum_{i=1}^{n_A} (z_j - y_{A,i})_+^{s-1} (z_k - y_{A,i})_+^{s-1} - \hat{D}_{A,j}^s \hat{D}_{A,k}^s \right] \frac{n}{n_A} \\ & + \left[\frac{1}{n_B [(s-1)!]^2} \sum_{i=1}^{n_B} (z_j - y_{B,i})_+^{s-1} (z_k - y_{B,i})_+^{s-1} - \hat{D}_{B,j}^s \hat{D}_{B,k}^s \right] \frac{n}{n_B} \\ & - \alpha_k \hat{g}_k [\hat{D}_{B,j}^s [1 - \hat{F}_B(y(r))] - \hat{D}_{A,j}^s [1 - \hat{F}_A(y(r))]] / \hat{f}(y(r)) \\ & - \alpha_j \hat{g}_j [\hat{D}_{B,k}^s [1 - \hat{F}_B(y(r))] - \hat{D}_{A,k}^s [1 - \hat{F}_A(y(r))]] / \hat{f}(y(r)) \\ & + \alpha_j \alpha_k \hat{g}_j \hat{g}_k [n_A \hat{F}_A(y(r)) [1 - \hat{F}_A(y(r))] \\ & + n_B \hat{F}_B(y(r)) [1 - \hat{F}_B(y(r))]] / [n \hat{f}^2(y(r))], \forall j, k, \end{aligned}$$

where $\hat{F}_A(y(r)) = n_A^{-1} \sum_{i=1}^{n_A} I(y_{A,i} \leq y(r))$, $\hat{g}_j = \hat{a}_j^B - \hat{a}_j^A + I(s = 1)[\hat{f}_B(\hat{z}_j) - \hat{f}_A(\hat{z}_j)]$,

$$\hat{a}_j^A = \frac{\partial \hat{D}_{A,j}^s}{\partial \hat{z}_j} = \frac{s-1}{n_A (s-1)!} \sum_{i=1}^{n_A} (\hat{z}_j - y_{A,i})_+^{s-2},$$

and \hat{f}_A is the estimated underlying density function of F_A .³ The estimates of distribution B are just the analogues of A .

To test the null hypothesis that distribution A stochastically dominates B at order s ,

$$H_0 : \Delta^s \geq 0.$$

The alternate hypothesis is simply the negation of H_0 .

Since we are testing multiple inequality restrictions, the relevant statistical inference methods can be found in (Kodde and Palm 1986). (Davidson and Duclos 2000) also uses this framework for their hypothesis tests. First, we compute the Wald-type test statistic

$$W = \min_{\Delta^s \geq 0} n(\hat{\Delta}^s - \Delta^s)' \hat{\Sigma}^{-1} (\hat{\Delta}^s - \Delta^s) \tag{6}$$

where the right-hand side is a quadratic programming problem. Under the null, W will converge in distribution to a mixture of χ^2 distributions.

³ A well-known method for estimating densities is kernel estimation. The consistency of such estimators has been rigorously established in the literature (see, e.g., (Li and Racine 2007)).

Obtaining critical values is not a straightforward process, so we follow (Davidson and Duclos 2000) and advocate the use of the bootstrap. The procedure can be explained as follows. Given samples y_A and y_B , we pool them and obtain $y = (y^A, y^B)'$. Then, the bootstrap samples y_A^* and y_B^* are generated by resampling n_A and n_B observations (with replacement) from y . Next, using the bootstrap samples, we compute the bootstrap test statistic W^* in a similar manner to W . After repeating this process, a large number of times, the bootstrap p -value is the proportion of times that W^* exceeds W . A value less than the nominal size of the test should lead to the rejection of H_0 .

Failure to infer dominance at order s by either distribution may imply that there exists some critical poverty line, $z_s < z$ where the distributions cross and thus a rank reversal occurs. If such a threshold exists, it can be implicitly characterized by $D_A^s(z_s) = D_B^s(z_s)$. A natural estimator of z_s is \hat{z}_s which solves $\hat{D}_A^s(\hat{z}_s) = \hat{D}_B^s(\hat{z}_s)$. Let $D^0(x) = f(x)$ and suppose $D_A^s(x) < D_B^s(x) \forall x < z_s$. (Davidson and Duclos 2000) showed in Theorem 3, in that case, $\sqrt{n_B}(\hat{z}_s - z_s)$ will be asymptotically normal with mean zero and asymptotic variance

$$\frac{\text{var}((z_s - y_A)_+^{s-1}) + n_B/n_A \text{var}((z_s - y_B)_+^{s-1})}{[(s-1)!(D_B^{s-1}(z_s) - D_A^{s-1}(z_s))]^2},$$

which, of course, can be estimated by simply replacing the terms with their respective empirical analogues.

4. Simulation Evidence

We now assess the size and power of our proposed test using a series of Monte Carlo experiments. The experiments were carried out using 10,000 independent trials, sample sizes of $n_A = n_B = n/2$, and nominal size set equal to 5%. We consider tests of first-order stochastic dominance (i.e., $s = 1$) for which $a_j^A = 0$. The poverty line is set equal to 50% of the pooled median (i.e., $c = q = 0.5$).

Five different parametric distributions are considered in assessing the size of the test: gamma, Singh–Maddala, log-normal, unit exponential, and uniform. The CDF of the gamma distribution is given by $F(y) = \gamma(a_2, y/a_1)/\Gamma(a_2)$, where a_1 is a scale parameter, a_2 is a shape parameter, $\gamma(\cdot)$ is the gamma function, and $\Gamma(\cdot)$ is the incomplete gamma function. The CDF of the Singh–Maddala distribution is given by $F(y) = 1 - (1 + b_1 y^{b_2})^{-b_3}$, where b_1 is a scale parameter, and b_2 and b_3 are shape parameters. Following (McDonald 1984), we set $a_2 = 2.1557$ for the gamma distribution, and $b_2 = 1.697$ and $b_3 = 8.368$ for the Singh–Maddala distribution, which were used to simulate 1980 U.S. income distribution. The scale parameters for both distributions are set to unity. For the log-normal distribution, the mean and standard deviation are set to 2.9372 and 0.7797, respectively, which were also used by (McDonald 1984) to simulate 1980 U.S. income distribution. For the uniform distribution, we follow (Zheng 2001) and specify the support as the unit interval $[0, 1]$. We consider five grid points set to 10%, 20%, 30%, 40%, and 50% of the pooled median.⁴

To assess the size of the test, we generate observations for subgroup A and B from the same distribution. We test the null hypothesis that $\Delta^1 \geq 0$, which is (weakly) true in this case. We consider pooled sample sizes varying from $n = 100$ to $n = 1000$ and utilize 199 bootstrap replications. Table 1 reports the rejection frequencies. Overall, we can conclude that a combined sample size of 1000 observations should be sufficient for achieving asymptotic normality. This is not a very demanding requirement at all as typical household survey datasets tend to have thousands of observations.

⁴ Since quantile-based poverty lines require density estimation in calculating the covariance structure, we use kernel estimation with a Gaussian kernel and a “rule-of-thumb” bandwidth (see, e.g., (Li and Racine 2007, Ch. 1)).

Table 1. Rejection frequencies for size simulation.

Distribution	<i>n</i>			
	100	300	500	1000
Gamma	6.32	5.07	5.02	4.68
Singh–Maddala	5.72	5.42	4.98	4.87
Log-normal	6.27	5.32	4.87	5.28
Unit Exponential	5.33	5.57	5.09	4.89
Uniform	5.20	5.28	4.62	4.88

Note: The nominal size of the test is 5%. The grid points are set to 10%, 20%, 30%, 40%, and 50% of the pooled median.

We focus exclusively on the gamma distribution in assessing the power of the test and consider the tests of (Davidson and Duclos 2000; Barrett and Donald 2003) as benchmarks.⁵ The shape and scale parameters for subgroup *A* remain set to their original levels from the size simulation. For subgroup *B*, we vary the parameters such that the CDF of *B* lies below the CDF of *A* for all points up to the poverty line. Rejection frequencies based on 199 bootstrap replications are reported in Table 2. The test exhibits excellent power properties as evidenced by the fact that it outperforms the test of (Barrett and Donald 2003) regardless of sample size. A direct comparison of our test with (Davidson and Duclos 2000) cannot really be made since their test is based on the assumption of either an absolute poverty line or poverty lines relative to each distribution (not pooled). Nonetheless, there is similarity between the two tests in terms of power. However, note that, in order to enable a somewhat fair comparison, the shape and scale parameters of the two distributions were chosen such that the medians remain the same for both distributions (thus, the pooled median is just the median of either subgroup *A* or subgroup *B*). This reduces sampling variability for the test of (Davidson and Duclos 2000) and permits a more valid comparison. The advantage offered by our proposed test is that it accounts for the sampling variance of the common poverty line that depends on both distributions while the other two tests do not.

Table 2. Rejection frequencies for power simulation.

Δ_1^1	Δ_2^1	Δ_3^1	Δ_4^1	Δ_5^1	<i>n</i>			
					100	300	500	1000
					<i>W</i>			
−0.01	−0.02	−0.04	−0.05	−0.05	23.22	44.47	61.84	86.92
−0.01	−0.04	−0.07	−0.09	−0.10	51.37	92.28	98.91	100.00
					<i>DD</i>			
−0.01	−0.02	−0.04	−0.05	−0.05	20.90	44.78	61.77	86.75
−0.01	−0.04	−0.07	−0.09	−0.10	52.07	92.32	98.83	100.00
					<i>BD</i>			
−0.01	−0.02	−0.04	−0.05	−0.05	13.63	24.10	33.24	51.09
−0.01	−0.04	−0.07	−0.09	−0.10	26.75	56.86	74.81	94.92

Note: The nominal size of the test is 5% and Δ_j^1 is the *j*th difference between the CDFs of subgroup *B* and *A* with grid points set to 10% (*j* = 1), 20% (*j* = 2), 30% (*j* = 3), 40% (*j* = 4), and 50% (*j* = 5) of the pooled median. Both distributions are generated from the gamma distribution. The scale parameter of distribution *A* is set to unity and its shape parameter is set to 2.1557. For distribution *B*, the scale is 0.651 and shape is 3.143 for the first row, while, for the second row, scale is 0.421 and shape is 4.688. DD denotes the test of (Davidson and Duclos 2000), and BD denotes the bootstrap test of (Barrett and Donald 2003).

⁵ For the case of a common relative poverty line, the test statistic of (Barrett and Donald 2003) is based on the supremum of the distances between the censored CDFs which sets all income values above the poverty line equal to the poverty line.

5. Illustration

In this section, we provide a simple example of how our proposed test can be applied in a real-world scenario. Using data from the 2017 Canadian Income Survey, we compare poverty outcomes among men and women. One of the ways Canada's national statistical agency, Statistics Canada, measures "low income" is through its low income measure (LIM), which sets the low income line equal to 50% of the median adjusted household income (household income is divided by the square root of the household size).⁶ The pooled median adjusted household income is determined to be \$46,461 implying a poverty line of \$23,230.50.

The benefit of using scalar poverty indices such as the low income measure is that it allows policy makers to monitor and assess trends across socio-economic groups, regions, and time. The downside is that robust comparisons are not assured since distributional comparisons are not made at points below the low income line. Since our proposed test assumes a common pooled relative poverty line such as Statistics Canada's low income line, we can make robust poverty comparisons between any subgroups of the population while still maintaining that common relative poverty line.

To illustrate our test, we compare poverty outcomes among Canadian men and women by testing for first-order stochastic dominance up to the low income line from the 2017 Canadian Income Survey. The sample consists of 47,800 men and 49,388 women. We consider five grid points set at 10%, 20%, 30%, 40%, and 50% of the median adjusted household income. The headcount ratios at the different points along with the median and standard deviation (SD) of adjusted household income for men and women are reported in Table 3. In testing the null that the male income distribution stochastically dominates the female income distribution up to the low income line, we obtain a bootstrap p -value of exactly zero leading to the rejection of the null. We also obtain a p -value of zero in testing the reverse hypothesis. This suggests that the CDFs of the two distributions cross at least once at some point below the low income line, which leads to our ambiguous conclusion that no dominance could be detected leading up to the low income line.

Table 3. Descriptive statistics and estimated headcount ratios from the 2017 Canadian Income Survey.

	Median	SD	\hat{D}_1^1	\hat{D}_2^1	\hat{D}_3^1	\hat{D}_4^1	\hat{D}_5^1
Men	\$47,581	\$52,766	0.008	0.016	0.034	0.062	0.113
Women	\$45,401	\$35,179	0.006	0.014	0.033	0.066	0.132

Note: D_j^1 denotes the headcount ratios at grid points set to 10% ($j = 1$), 20% ($j = 2$), 30% ($j = 3$), 40% ($j = 4$), and 50% ($j = 5$) of the pooled sample median (\$46,461).

From Table 3, observe that, according to the low income line, the poverty rate is 11.3% for men and 13.2% for women (note that our estimates vary slightly from the official estimates because we treat negative incomes as zero and ignore the complex sampling scheme of the survey to simplify this exposition).⁷ Notice that, at the first three grid points, the poverty rates are lower for women. The trend reverses at the higher grid points which suggests that the CDFs cross, negating the possibility of first-order dominance by either subgroup.

However, can we make any inference regarding the critical poverty line where the rank reversal occurs? In other words, we want to determine the threshold level of income up to which the women's distribution dominate the men's. We determine first-order stochastic dominance of women's income

⁶ The low-income cut-offs (LICOs) and market basket measures (MBMs) are two of the other complementary ways Statistics Canada measures and monitors low income and poverty. Unlike the LIM, the LICO and MBM are absolute poverty lines that vary by region and attempt to account for cost-of-living. As of 2018, the MBM became Canada's official poverty line. However, the LIM continues to be the most commonly used measure for international poverty comparisons. For detailed information regarding the different measures, see (Statistics Canada 2016).

⁷ (Chen and Duclos 2011) provide an illustration of stochastic dominance tests using Canadian data that take into account sampling weights.

distribution over men up to an adjusted household income level of \$15,058, which is 32% of the median. The 95% confidence interval for the critical threshold in our case is \$12,731.56 to \$17,384.44, or 27% to 37% of the median.

6. Conclusions

In this article, we proposed a test for stochastic dominance up to a common relative poverty line. Much of the existing tests of stochastic dominance, in the context of poverty measures, assume either separate poverty lines for each distribution or a common absolute poverty line. It is increasingly the case that relative poverty lines (i.e., 50% of the median income level) are being used in cross-country and group comparison studies.

A series of Monte Carlo experiments validates our asymptotic framework. The proposed test exhibits good size and power properties, under varying conditions, and outperforms existing methods due to the fact that this test utilizes the underlying covariances between estimates at the different points of the distributions. A sample size of 1,000 observations appears to be sufficient for achieving asymptotic normality which is not a very demanding requirement as household surveys tend to have thousands of observations. For illustration purposes, household income data from the 2017 Canadian Income Survey were used to rank poverty outcomes among men and women using the Canadian national statistical agency's low income line.

Funding: This research received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

References

- Anderson, Gordon 1996. Nonparametric tests of stochastic dominance in income distributions. *Econometrica* 64: 1183–93.
- Atkinson, Anthony Barnes 1987. On the measurement of poverty. *Econometrica* 55: 749–64.
- Barrett, Garry F., and Stephen G. Donald. 2003. Consistent tests for stochastic dominance. *Econometrica* 71: 71–104.
- Barrett, Garry F., Stephen G. Donald, and Yu-Chin Hsu. 2016. Consistent tests for poverty dominance relations. *Journal of Econometrics* 191: 360–73.
- Chen, Wen-Hao, and Jean-Yves Duclos. 2011. Testing for poverty dominance: An application to Canada. *Canadian Journal of Economics* 44: 781–803.
- Davidson, Russell, and Jean-Yves Duclos. 2000. Statistical inference for stochastic dominance and the measurement of poverty and inequality. *Econometrica* 68: 1435–64.
- Foster, James, Joel Greer, and Erik Thorbecke. 1984. A class of decomposable poverty measures. *Econometrica* 52: 761–66.
- Foster, James 1984. On economic poverty: A survey of aggregate measures. *Advances in Econometrics* 3: 215–51.
- Garcia-Gomez, Cesar, Ana Perez, and Mercedes Prieto-Alaiz. 2019. A review of stochastic dominance methods for poverty analysis. *Journal of Economic Surveys* 33: 1437–62.
- Kakwani, Nanak 1993. Statistical inference in the measurement of poverty. *The Review of Economics and Statistics* 75: 632–39.
- Kaur, Amarjot, Bhagavatula Lakshmi Surya Prakasa Rao, and Harshinder Singh. 1994. Testing for second-order stochastic dominance of two distributions. *Econometric Theory* 10: 849–66.
- Kodde, David A., and Franz C. Palm. 1986. Wald criteria for jointly testing equality and inequality restrictions. *Econometrica* 54: 1243–48.
- Li, Qi, and Jeffrey Scott Racine. 2007. *Nonparametric Econometrics: Theory and Practice*. Princeton: Princeton University Press.
- McDonald, James B. 1984. Some generalized functions for the size distribution of income. *Econometrica* 52: 647–63.
- McFadden, Daniel 1989. Testing for stochastic dominance. In *Studies in the Economics of Uncertainty*. Edited by T. B. Fomby and T. K. Seo. New York: Springer.

- Mehdi, Tahsin 2017. Poverty comparisons with common relative poverty lines. *Communications in Statistics—Theory and Methods* 46: 2029–36.
- OECD. 2016. Poverty rates and gaps. In *OECD Factbook 2015-2016: Economic, Environmental and Social Statistics*. Paris: OECD Publishing. doi:10.1787/factbook-2015-21-en.
- Ravallion, Martin 1994. *Poverty Comparisons, Fundamentals of Pure and Applied Economics*. Chur, Langhorne: Harwood Academic Publishers.
- Sen, Amartya 1976. Poverty: An ordinal approach to measurement. *Econometrica* 44: 219–31.
- Statistics Canada. 2016. *Low Income Lines: What They Are and How They Are Created*. Ottawa: Statistics Canada.
- Thompson, Brennan Scott, and Thanasis Stengos. 2012. Testing for bivariate stochastic dominance using inequality restrictions. *Economics Letters* 115: 60–62.
- Thompson, Brennan Scott 2012. Empirical likelihood-based inference for poverty measures with relative poverty lines. *Econometric Reviews* 32: 513–23.
- Zheng, Buhong 2001. Statistical inference for poverty measures with relative poverty lines. *Journal of Econometrics* 101: 337–56.



© 2020 by the author. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).