Contributions to the measurement of relative p-bipolarisation[1]

Marek Kosny
University of Wroclaw

and

Gaston Yalonetzky
University of Leeds

This version: January 2017

Abstract

Firstly, we axiomatically characterize the class of relative bipolarisation indices. Then we propose the first class of relative bipolarisation measures which are simultaneously *percentile-independent and partially rank-independent*. These are based on differences of generalised means. We also propose a relative bipolarisation pre-ordering based on pairs of *hybrid Lorenz curves* which combine features of both relative and generalised Lorenz curves. Considering different ways to divide distributions into two mutually exclusive and exhaustive groups using a percentile (e.g. the median), we also characterize the very few instances in which one distribution can dominate another one in terms of relative bipolarisation across the whole percentile domain. We illustrate the measures and curves with a comparison of the US versus Germany across time.

1. Introduction

Bipolarisation indices and pre-orderings have gained traction as methods to measure the growth or disappearance of middle-classes for the last few decades, since the foundational work of Foster and Wolfson (2010; based on a 1992 working paper) and Wolfson (1994). Essentially, in the original conception by Foster and Wolfson, bipolarisation measurement requires partitioning distributions into two groups using a dividing percentile (usually the median), and then distinguishing between transfers across that percentile and transfers on one side of that percentile (i.e. within one group). As with inequality measurement, a regressive transfer across the dividing percentile is deemed to increase the spread of mean attainment between the two groups, thereby increasing bipolarisation. By contrast, unlike inequality measurement, a progressive transfer within any one group is deemed to increase clustering, in the limit leading to perfect bimodality; hence these progressive transfers are deemed to increase bipolarisation.

Above and beyond the common treatment of transfers, bipolarisation indices differ in many ways; a key one being their reaction to changes in the unit of income measurement. Thus,

[1] The authors would like to thank participants at the 34th IARIW General Conference, Dresden, August 2016, for their helpful comments.

there are relative, absolute, and intermediate classes of indices.[2] In this paper we focus on relative bipolarisation indices and pre-orderings, i.e. those fulfilling a property of *(ratio-)scale invariance*, whereby changes in the unit of income measurement do not alter the value of the indices. Admittedly one could opt for a less stringent property of unit-consistency (e.g. see Lasso de la Vega, 2010; for bipolarisation measurement), which only requires the ranks in any pair-wise comparison to remain unaltered when the unit of income measurement changes. However, focusing on a relative approach bears the advantage of working with clear benchmarks of both minimum and maximum bipolarisation, in addition to being the most straightforward framework to introduce our contributions, which can be generalised to alternative bipolarisation measurement approaches.

We provide four methodological contributions to the measurement of relative bipolarisation. Firstly, we axiomatically characterize the class of relative bipolarisation indices. Our result is an application of Bossert and Schworm (2008, Theorem 3) to classes of indices expected to satisfy the property of (ratio-)scale invariance defining the relative approach.[3]

Secondly, we introduce the first class of relative bipolarisation indices which are simultaneously median-independent (in fact, percentile-independent) and partially rank-independent. These indices are based on normalised differences of generalised means. Relative bipolarisation indices can be classified into median-dependent or median-independent (in general, we could say percentile-dependent or percentile-independent). Examples of the former include the famous Foster-Wolfson index, but also the class $P_2^N(x)$ by Wang and Tsui (2000, p. 358). Examples of the latter include proposals by Wang and Tsui (2000), e.g. the class $P_1^N(x)$, and by Yalonetzky (2016b). Unfortunately, as shown by Yalonetzky (2016b), median-dependent indices violate the key transfer axioms of bipolarisation, unless the median remains unaltered by transfers, which is both unnecessarily restrictive and not guaranteed in practice. Hence we can effectively rely only on median-independent indices. To date, all median-independent indices of relative bipolarisation proposed in the literature are rank-dependent, even though we know, from Bossert and Schworm (2008, Theorem 2), that full rank-dependence is not a required feature in the construction of sound bipolarisation indices. In this context, we propose the first class of median-independent, partially rank-independent indices, which also provide the computational advantage of rendering any rank function superfluous. The only rank-related requirement is partitioning the population into two non-overlapping groups (hence the partial nature of rank independence). Additionally, we show that the class of indices is easily decomposable into a spread component and a clustering component.[4]

---

[2] Examples of relative indices include those by Foster and Wolfson (2010), Wang and Tsui (2000), and Deutsch et al. (2007). Examples of absolute indices include the general class by Bossert and Schworm (2008). Finally, examples of intermediate indices include the family by Chakravarty and D'Ambrosio (2010).

[3] Hereafter we refer to ratio-scale invariance as simply scale invariance.

[4] There exist fully rank-independent indices in the literature, e.g. classes $P_3^N(x)$ and $P_4^N(x)$ by Wang and Tsui (2000, p. 359) but they are all median-dependent, rendering them unsuitable for comparisons when medians can be affected by transfers (Yalonetzky, 2016b).

Thirdly, inspired by the seminal paper of Bossert and Schworm (2008), we derive a pre-ordering for relative bipolarisation measurement based on *hybrid Lorenz curves* which combine features of both relative and generalised Lorenz curves. So far, the literature has provided two proposals for pre-orderings of relative bipolarisation, specifically. The original proposal by Foster and Wolfson (2010) was based on their second-order bipolarisation curves, which are median-dependent, hence unfortunately not suitable for its intended purpose, due to the same reasons put forward by Yalonetzky (2016b) in order to dismiss median-dependent indices. The second proposal, by Yalonetzky (2014), is based on so-called relative bipolarisation Lorenz curves. These are median-independent, therefore suitable for relative bipolarisation pre-ordering. However, it is not easy to adapt them for alternative partitions of the distribution, i.e. using a percentile different from the median. By contrast, our proposed pre-orderings based on hybrid Lorenz curves are more flexible: they are suitable for relative bipolarisation measurement with any choice of dividing percentile.

Fourthly, following the concluding reflections by Bossert and Schworm (2008), we first emphasise how the above contributions (new class of indices and new pre-ordering), like some previous ones in the literature, are applicable to any partition of distributions into two groups (i.e. not just identical halves using the median). Thus we formally introduce the concept of relative *p-bipolarisation*. For instance, in addition to naturally choosing the median and using bipolarisation measurement to gauge the rise or demise of the middle class, we could also choose the 95^{th} percentile and test whether the top 5% is getting more clustered while separating itself from the rest of society. Lastly, and relying on the hybrid Lorenz curves, we contribute by characterizing the very few situations in which one distribution can dominate another one in terms of relative bipolarisation *across the whole percentile domain*. Hence, unless societies A and B correspond to any of these exceptional situations, A can never dominate B over every possible partition of society into two mutually exclusive and exhaustive groups, and vice versa.

We illustrate the usefulness of our indices, and of measuring bipolarisation using different dividing percentiles, with a comparison of household and individual income between the United States and Germany. Interestingly, we find that relative bipolarisation is higher for individuals in the United States (pre-government income), but the situation is reversed for households (also pre-government income). However, resorting to the hybrid Lorenz curves, we show that the results were not robust to any choice of relative bipolarisation index.

The observed higher bipolarisation of household income in Germany occurred despite large income inequalities at the top of the income distribution in the United States (higher than in Germany). It was mainly led by the significantly higher percentage of households without any income in Germany. Additional analysis of the differences between pre-and post-government income suggests that an important factor explaining the observed difference in household income bipolarisation can be the institutions of the welfare state. In Germany, they guarantee an acceptable standard of living also in a situation where the household receives no market income. Another phenomenon observed in the United States, and absent in Germany, is the relative impoverishment of a group of people with incomes above the median, but not belonging to the top 5%.

The rest of the paper proceeds as follows. Section 2 provides the notation and the definition of key statistics, benchmarks of minimum and maximum relative bipolarisation, followed by the main axioms. Section 3 provides the axiomatic characterization of the class of relative bipolarisation indices. Section 4 introduces our class of percentile-independent and rank-independent indices of relative bipolarisation, showing that it satisfies all the key desirable axioms, and that it is easily decomposable into a spread component and a clustering component. Section 5 develops the pre-orderings for relative p-bipolarisation based on hybrid Lorenz curves. Section 6 is dedicated to show that relative p-bipolarisation dominance cannot exist over the whole percentile domain, save for two types of distributional comparisons. Section 7 provides the empirical illustration. Then the paper concludes with some final remarks.

2. Notation and preliminaries

Let $y_i \geq 0$ denote the income of individual i. Y is the income distribution with mean $\mu_Y > 0$, and size $N \geq 4$.[5] Individuals are ranked in ascending order so that: $y_1 \leq y_2 \leq \cdots \leq y_{N-1} \leq y_N$. We denote a percentile with $p \in [0,1] \subset \mathbb{R}_+$. We will also be using quantile functions of the form $y(p)$, such that, for instance, $y(0.5)$ is the median of Y. For any possible p we also define the bottom part of Y as $\underline{Y}(p) = \{y_i \in Y: y_i \leq Y(p)\}$, as well as the top part $\overline{Y}(p) = \{y_i \in Y: y_i > Y(p)\}$,

We further define a bipolarisation index $B: Y \to \mathbb{R}_+$. It will also be useful to define a rank-preserving *Pigou-Dalton progressive transfer*, involving incomes $y_i < y_j$ and a positive amount $\delta > 0$ such that: $y_i + \delta \leq y_j - \delta$. If the transfer between the same pair is in the opposite direction, i.e. favouring the already wealthier individual, then we call it a *regressive transfer*.

Following, Yalonetzky (2016a) we should also define two sets of distributions which provide the benchmarks of minimum and maximum relative bipolarisation. The first set, $\mathcal{E}$, is made of distributions exhibiting equal non-negative incomes. That is $\mathcal{E} = \{Y \in \mathbb{R}^N_{++}: y_1 = y_2 = \cdots = y_N = y > 0\}$. This is the set of all perfectly egalitarian distributions, which the literature also defines as the benchmark of minimum bipolarisation. The second set, $\mathcal{B}_p$, is made of a bottom p of null incomes and a top $1-p$ of egalitarian incomes. That is $\mathcal{B}_p = \{Y \in \mathbb{R}^N_+: y_1 = y_2 = \cdots = y_{pN} = 0 \wedge y_{pN+1} = y_{pN+2} = \cdots = y_N = y > 0\}$. This is the set that Yalonetzky (2016a)

[5] For the measurement of bipolarisation, ideally we would like to have at least two people on each of the two parts of the distribution.

characterises as the benchmark of maximum bipolarisation given a partition of the population into two adjacent non-overlapping parts: the bottom p and the top $1-p$.[6]

Now we define the generalised means of the bottom and top parts:

$$\underline{\mu}(Y;p,\alpha) \equiv \left[\frac{1}{Np}\sum_{i=1}^{Np} y_i^{\alpha}\right]^{\frac{1}{\alpha}}, \forall\alpha \neq 0 \quad (1)$$

$$\overline{\mu}(Y;p,\beta) \equiv \left[\frac{1}{N(1-p)}\sum_{i=Np+1}^{N} y_i^{\beta}\right]^{\frac{1}{\beta}} \; \forall\beta \neq 0 \quad (2)$$

$\underline{\mu}(Y;p,\alpha)$ is the generalised mean of the bottom part and $\overline{\mu}(Y;p,\beta)$ is the generalised mean of the top part.

The next step is to list the desirable properties for an index of relative p-bipolarisation. We start with axioms 1 and 2 which are standard in the literatures on inequality, polarisation and bipolarisation:

Axiom 1: Symmetry (SY): $B(X;p) = B(Y;p)$ if $X = VY$ where V is an $N \times N$ permutation matrix.

Axiom 2: Population principle (PP): $B(X;p) = B(Y;p)$ if $X \in \mathbb{R}_+^{\lambda N}$ is obtained from $Y \in \mathbb{R}_+^N$ through an equal replication of each individual income, λ times.

The axiom that distinguishes the relative approach to bipolarisation (and inequality measurement) from other unit-consistent approaches (e.g. absolute, intermediate) is scale invariance:

Axiom 3: Scale invariance (SC): $B(X;p) = B(Y;p)$ if $X = \theta Y$ and $\theta > 0$.

Then we present the two classic transfer axioms of bipolarisation (Foster and Wolfson, 2010). Axiom 4 states that a Pigou-Dalton transfer involving incomes from the bottom and the top part should decrease the value of the bipolarisation index, as the *spread* between the two parts is narrowed down. By contrast, axiom 5 requires an increase in the value of the bipolarisation index, whenever a Pigou-Dalton transfer takes place between either two incomes of the bottom part or two incomes of the top part, since any of the latter implies an increase in the degree of clustering within the parts.

Axiom 4: Spread-decreasing Pigou-Dalton transfers (SD): If X is obtained from Y through PD transfers *across the* $y(p)$ *quantile*, which do not make any affected income switch the part of the distribution (bottom or top) to which they initially belonged, then $B(X;p) < B(Y;p)$.

Axiom 5: Clustering-increasing Pigou-Dalton transfers (CI): If X is obtained from Y through PD transfers *on one side of the* $y(p)$ *quantile* then $B(X;p) > B(Y;p)$.

[6] Yalonetzky (2016) does this characterisation for $p = 0.5$, but it is easy to show that $\mathcal{B}_p$ provides the benchmark of maximum relative bipolarisation for any chosen p.

It will be useful later to consider also the equivalent counterparts of axioms 4 and 5 defined in terms of regressive transfers:

Axiom 4a: Spread-increasing regressive transfers (SR): If X is obtained from Y through regressive transfers *across the* $y(p)$ *quantile* then $B(X;p) > B(Y;p)$.

Axiom 5a: Clustering-decreasing regressive transfers (CR): If X is obtained from Y through regressive transfers *on one side of the* $y(p)$ *quantile,* which do not make any affected income switch the part of the distribution (bottom or top) to which they initially belonged, then $B(X;p) < B(Y;p)$.

Finally, we include the normalisation axiom for relative bipolarisation measurement, which is the only one consistent with previous axioms (as shown by Yalonetzky, 2016a):

Axiom 6: Normalisation (N): (a) $B(Y;p) > B(X;p) = 0$ if and only if $X \in \mathcal{E}$ and $Y \notin \mathcal{E}$, and (b): $B(Y;p) < B(X;p) = 1$ if and only if $X \in \mathcal{B}_p$ and $Y \notin \mathcal{B}_p$.

3. The class of relative bipolarisation indices

Both for its own sake and as a preliminary step to proposing an appealing new specific class of relative bipolarisation indices (in the next section), we deem it worthwhile to characterize the general class of relative bipolarisation indices satisfying the two key transfer axioms and the defining property of scale invariance. For reasons presented by Yalonetzky (2016b) the median (or any other quantile) must be absent from the index's formula. Theorem 1 characterizes the class of relative bipolarisation indices:

Theorem 1: $B(Y;p)$ satisfies axioms (SR), (CI), and (SC) if and only if (a) $B(Y;p)$ is of the form $B(Y;p) = f(\frac{\overline{Y}}{N\mu_Y}, \frac{\underline{Y}}{N\mu_Y})$, (b) $B(Y;p)$ is increasing and Schur-concave in $\overline{Y}$, and (c) $B(Y;p)$ is decreasing and Schur-concave in $\underline{Y}$.[7]

Proof: See Appendix.

Theorem 1 is clearly inspired by Bossert and Schworm (2008, Theorem 2), especially aspects (b) and (c). However, two differences are worth pointing out. Firstly, the requirement that $B(Y;p)$ be a function of $\frac{\overline{Y}}{N\mu_Y}$ and $\frac{\underline{Y}}{N\mu_Y}$, as opposed to $\overline{Y}$ and $\underline{Y}$ (Bossert and Schworm, 2008, p. 1179). Secondly, we define the spread property in terms of a regressive transfer (following Foster and Wolfson, 2010, in turn based on their 1992 working paper), whereas Bossert and Schworm (2008, p. 1174) define it in terms of vector dominance (e.g. if $\underline{Y} \leq \underline{X}$ and $\overline{Y} \geq \overline{X}$ then Y is deemed more bipolarised than X). While the two definitions may be identical in many situations, the transfer definition always revolves around mean-preserving transfers. By contrast, the vector dominance definition does not require these transfers. As a starting point,

[7] $N\mu_Y = y_1 + y_2 + \cdots + y_N$. That is, theorem 1 states that $B(Y;p)$ must be a function of *income shares* in order to satisfy the three key axioms of relative bipolarisation measurement.

we find it easier to reconcile scale invariance with the spread property when the latter is defined in terms of mean-preserving transfers.[8]

4. Quantile-independent, rank-independent indices of relative bipolarisation based on generalised means

4.1. Definition

Consider the following functional form for a quantile-independent and partially rank-independent index of relative bipolarisation:

$$B(Y;p,\alpha,\beta) \equiv \frac{1-p}{\mu_Y}\left[\overline{\mu}(Y;p,\beta) - \underline{\mu}(Y;p,\alpha)\right] \quad (3)$$

Clearly, $B(Y;p,\alpha,\beta)$ already fulfils symmetry, population principle, and scale invariance.

Now we characterize the subset (α,β) which renders $B(Y;p,\alpha,\beta)$ a suitable class of relative bipolarisation indices:

Proposition 1: $B(Y;p,\alpha,\beta)$ fulfils SD, CI, and N if and only if $\alpha > 1 > \beta$.

Proof: See Appendix.

4.2. Decomposition

An appealing trait of $B(Y;p,\alpha,\beta)$ with $\alpha > 1 > \beta$ is that it is easily decomposable into a spread component and a clustering component. That is, if we track relative p-bipolarisation across time or compare two countries, regions, etc., we can compute the proportion of the difference in bipolarisation which is due to greater spread between the means of the two parts of the distribution, and the proportion which is due to differential clustering within each part.

The decomposition works as follows: Firstly, note that $B(Y;p,1{,}1)$ is simply the mean-normalized difference between the mean of the top part and the mean of the bottom part. It is straightforward to verify that $B(Y;p,1{,}1)$ fulfills N and SD, but not CI. That is, $B(Y;p,1{,}1)$ is insensitive to any type of transfers within either of the parts. Secondly, note that $B(Y;p,\alpha,\beta)$ with $\alpha > 1 > \beta$ can be decomposed into two components, one of which is $B(Y;p,1{,}1)$:

$$B(Y;p,\alpha,\beta) = [B(Y;p,\alpha,\beta) - B(Y;p,1{,}1)] + B(Y;p,1{,}1) \quad (4)$$

Thirdly, let $C(Y;p,\alpha,\beta) \equiv B(Y;p,\alpha,\beta) - B(Y;p,1{,}1)$, and note that $C(Y;p,\alpha,\beta)$ fulfills CI, as long as $\alpha > 1 > \beta$. Hence we can attribute the clustering effect to $C(Y;p,\alpha,\beta)$ while measuring the spread effect with $B(Y;p,1{,}1)$. Also note that with $\alpha > 1 > \beta$, we have $C(Y;p,\alpha,\beta) \leq 0$, and $C(Y;p,\alpha,\beta) = 0$ only in the absence of inequality within each and

[8] There is also a third difference: since our transfer axioms are defined in terms of strict inequalities, we can state, in theorem 1, that $B(Y;p)$ is strictly increasing in $\overline{Y}$ and decreasing in $\underline{Y}$; whereas theorem 2 by Bossert and Schworm (2008) refers to non-decreasing and non-increasing behaviour, respectively.

every part of the distribution. This means that an increase in clustering leads to a higher $B(Y; p, \alpha, \beta)$ through a lower absolute value of $C(Y; p, \alpha, \beta)$. Finally, note that the choice of (α, β) affects the relative size of the two effects. For any given degree of inequality within both parts, the relative importance of the clustering effect diminishes as both parameters, α and β, tend toward 1.

5. Pre-orderings for relative p-bipolarisation

As mentioned above, the two existing proposals for pre-orderings of *relative* bipolarisation suffer from limitations that warrant a new development. The pre-ordering of Foster and Wolfson (2010) is, in principle, adjustable to percentile partitions different from the median, but it is a function of the quantile itself, which leads to the violation of the transfer axioms, as shown by Yalonetzky (2016b). Meanwhile, the pre-ordering of Yalonetzky (2014) does not depend on the median, but is difficult to generalise to uneven partitions of the population into two groups.

Bossert and Schworm (2008) proposed a pre-ordering for bipolarisation based on generalised Lorenz curves. This pre-ordering is indeed flexible to any percentile partition. However, as it stands, it is not consistent with scale invariance; therefore, it is more suitable for an absolute conception of bipolarisation. For instance, if one compared the UK income distribution in pound sterling versus the same distribution measured in US dollars using this pre-ordering, then one would conclude that the pound-denominated distribution is robustly less bipolarised than the dollar-denominated distribution, for any percentile partition, simply because the absolute spread was widened by the exchange rate.

Hence, if we want to perform robust comparisons of relative p-bipolarisation, i.e. with flexibility in the choice of dividing percentiles, we need a suitable pre-ordering. We build on the seminal idea of Bossert and Schworm (2008), but instead of generalised Lorenz curves we construct our proposal relying on so-called *hybrid Lorenz curves*.

Hybrid Lorenz curves are essentially relative Lorenz curves that accumulate incomes, ordered from lowest to highest, by mapping only from a convex subset of the income distribution domain. Just like relative Lorenz curves, they have the mean in the denominator, thereby being consistent with the axiom of scale invariance. However, unlike relative Lorenz curves and rather akin to generalised Lorenz curves, the highest value of hybrid Lorenz curves is variable and equal to the ratio between the mean of the chosen subset of incomes and the total population (or sample) mean. This ratio can take values in the subset $[1,2] \subset \mathbb{R}_{++}$ when $p = 0.5$. More generally, the ratio will be bounded by the subset $\left[1, \frac{1}{1-p}\right] \subset \mathbb{R}_{++}$ Hence the hybrid nature of the curves.

Formally, we define the hybrid Lorenz curve as a mapping function $L(Y, q, r; k)$: $\mathbb{R}_{+}^{N+2} \rightarrow \mathbb{R}_{+}$, where $q < r$ are percentiles and $k = qN, qN + 1, \dots, rN$:

$$L(Y, q, r; k) \equiv \frac{1}{rN - qN + 1} \sum_{i=qN}^{k} \frac{y_i}{\mu_Y} , \quad k = qN, qN + 1, \dots, rN \qquad (5)$$

It will also be helpful to define a *reverse hybrid Lorenz curve*, which accumulates incomes from a convex subset of the income domain, ordered from highest to lowest, and whose maximum value is also variable and equal to the ratio between the mean of the chosen subset of incomes and the total population (or sample) mean (taking values in the subset $[0,1] \subset \mathbb{R}_{++}$):

$$RL(Y,q,r;k) \equiv \frac{1}{rN-qN+1} \sum_{i=k}^{rN} \frac{y_{rN-i+1}}{\mu_Y},, \quad k = qN, qN+1, \dots, rN \quad (6)$$

Now in the particular case of relative bipolarisation measurement, once we select a percentile p and use it to divide the distribution into two parts, we can define two hybrid Lorenz curves, one for each of these two parts:

$$L(Y,p,1;k) \equiv \frac{1}{N-Np} \sum_{i=pN+1}^{k} \frac{y_i}{\mu_Y}, \quad k = pN+1, pN+2, \dots, N \quad (7)$$

$$RL(Y,0,p;k) \equiv \frac{1}{Np} \sum_{i=k}^{pN} \frac{y_{pN-i+1}}{\mu_Y}, \quad k = 1,2, \dots, pN \quad (8)$$

Figure 1 shows some of the possible shapes that $L(Y,p,1;k)$ and $RL(Y,0,p;k)$ can take under different situations of relative bipolarisation, in the specific case of $p = 0.5$.[9] For example $L(E,0.5,1;k)$ and $RL(E,0,0.5;k)$, which appear in green, are basically the hybrid Lorenz curves of any egalitarian distribution, i.e. $E \in \mathcal{E}$. Meanwhile, $L(M,0.5,1;k)$ and $RL(M,0,0.5;k)$, both in red, correspond to the hybrid Lorenz curves to a distribution characterized by maximum relative bipolarisation, i.e. $M \in \mathcal{B}_p$. The curves $L(A,0.5,1;k)$ and $RL(A,0,p;k)$, both in brown, depict a situation of perfect bimodality (in which the bottom part features positive incomes), but short of maximum relative bipolarisation (in which the bottom part features only null incomes). Finally, $L(B,0.5,1;k)$ and $RL(B,0,0.5;k)$, both in blue, are the hybrid Lorenz curves of a more typical distribution characterized by some degree of inequality both between its two parts (spread) and within each of them (imperfect clustering or lack of perfect bimodality).

Figure 1: Examples of hybrid Lorenz curves

[9] Note the different origins on the intersection between the two axes in Figure 1. Also note the following feature (which is relevant in the proofs below): $L(Y,p,1;1) = L(X,p,1;1) \leftrightarrow RL(Y,0,p;p) = RL(X,0,p;p)$. The proof is straightforward upon realising that $L(Y,0.5,1;1) + RL(Y,0,0.5;0.5) = 2$.

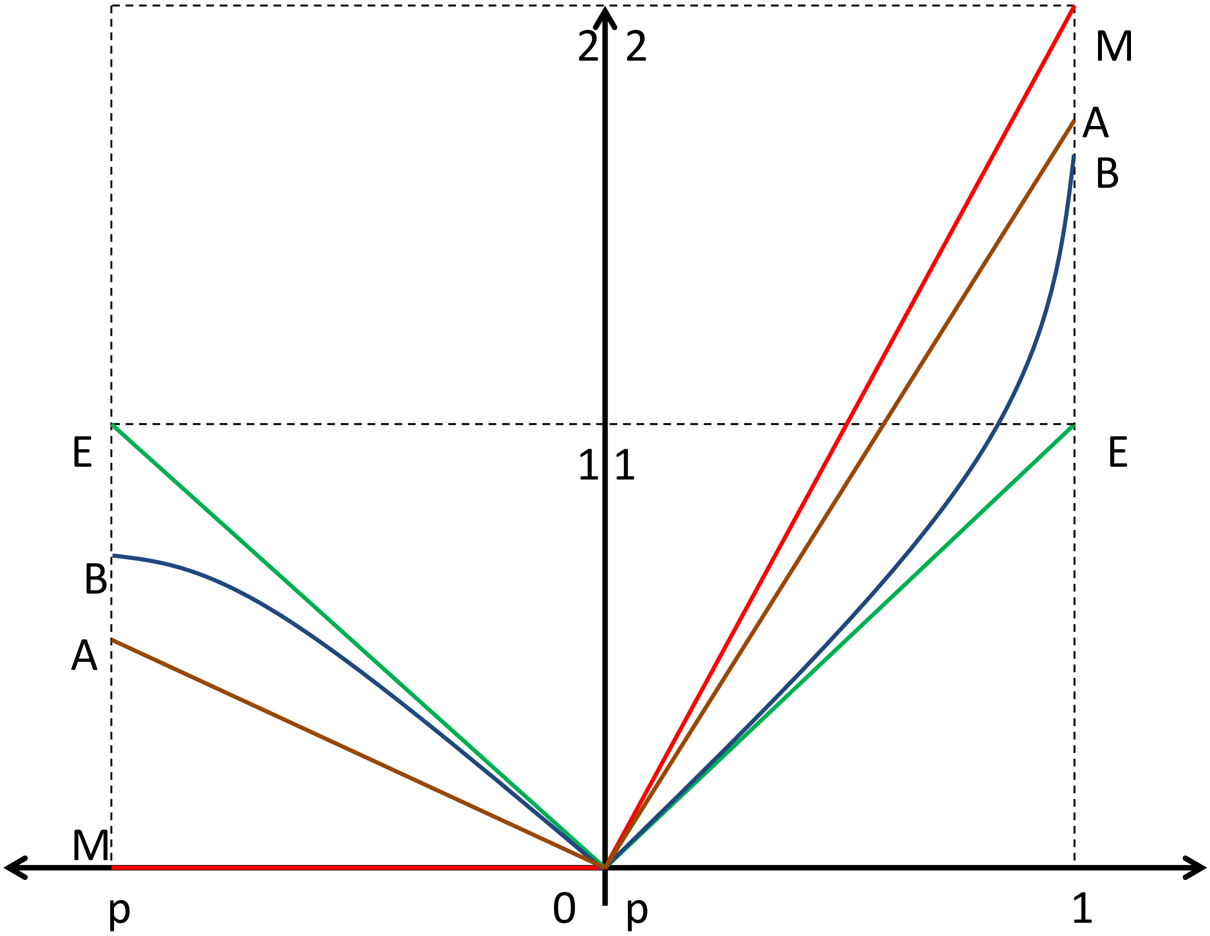

Now we can state the theorem that guarantees robust relative bipolarisation comparisons for a given distributional partition based on percentile p:

Theorem 2: $B(X;p) > B(Y;p)$ for all B satisfying SY, PP, SC, SD/SR, CI/CR, and N if and only if (i) $L(X,p,1,;k) \geq L(Y,p,1;k)\ \forall k = pN+1, pN+2, \ldots, N$, with at least one strict inequality; and (ii) $RL(X,0,p;k) \leq RL(Y,0,p;k)\ \forall k = 1,2,\ldots,pN$, with at least one strict inequality.

Proof: See Appendix.

Essentially Theorem 1 requires the hybrid Lorenz curve of the top part of X to be never below Y's, and at least once strictly above, while requiring the reverse hybrid Lorenz curve of the bottom part of X to be never above Y's, and at least once strictly below. In the case of the exemplary distributions in Figure 1, then Theorem 1 states that distribution A exhibits robustly more relative bipolarisation than B, and both are robustly more bipolarised than the egalitarian distribution E. Finally, distribution M is robustly more bipolarised than all the others, given that, in fact, it represents the benchmark of maximum relative bipolarisation when the partition is based on percentile p.

Finally note that Theorem 1 of Bossert and Schworm (2008) coincides with Theorem 1 above, only when both the means and population sizes are identical.

6. Limits to relative p-bipolarisation dominance

Bossert and Schworm (2008) mentioned that their framework could be generalised to any partition of the distribution into two groups, i.e. not just a partition into equally sized parts. Of course, for a bipolarisation assessment to make sense we should at least require: $pN \geq 2$ and $(1-p)N \geq 2$. Hence in addition to making the traditional bipolarisation comparisons using $p = 0.5$, we could also, for instance, apply these tools to assess whether the "top 5%" and the "rest" are clustering within while spreading away from each other. In that case we would choose $p = 0.95$.

In this section we explore further questions that arise when we count on several choices for p. In particular we pose two questions: (1) Consider distributions $A, B \notin \mathcal{E}$, if A dominates B for a given percentile, can it dominate B for any other percentile? (2) consider now two more general distributions A and B, can A dominate B throughout the whole relevant percentile domain, i.e. $]\frac{2}{N}, 1-\frac{2}{N}[$?

We start with the first question: can one non-egalitarian distribution dominate another non-egalitarian distribution across more than one percentile? The answer is a direct "yes". Consider for example the following two distributions: $A = \{1,2,3,4,5,6,7,8,9,10\}$ and $B = \{0,2.5,2.5,4,5,6,7,8.5,8.5,11\}$. Clearly, B was obtained from A through one regressive transfer (involving 1 and 10) and two Pigou-Dalton transfers (involving the pairs 2 and 3, and 8 and 9). If we choose $p = 0.5$ we will find that B is robustly more bipolarised than A. But we will reach the same conclusion if we choose, alternatively, $p = 0.4$ or $p = 0.6$. For the three choices, the Pigou-Dalton transfers take place between incomes on the same side of the partition, whereas the regressive transfer occurs across all partitions.

The second question follows naturally: Could B actually dominate A over the whole relevant percentile domain? Here the answer is slightly less straightforward, and we state it as a theorem:

Theorem 3: (i) Distribution A cannot dominate B over the whole relevant percentile domain, unless: (ii) $A \notin \mathcal{E} \wedge B \in \mathcal{E}$; and (iii) A was obtained from B through a sequence of regressive transfers involving, each time, one income in percentile $q \leq \frac{2}{N}$ and one income in percentile $r \geq 1-\frac{2}{N}$.

Proof: See Appendix.

Essentially, Theorem 3 states that, in general, dominance cannot be established over the whole relevant percentile domain, even though it can be established over a portion of the domain (as we saw in the answer to the first question). However, there are only two narrow cases in which dominance holds over the whole domain: (1) When one distribution is egalitarian and the other one is not; (2) When one distribution is obtained from the other one through regressive transfers involving the two lowest incomes and the two highest incomes.

7. Analysis of relative bipolarisation in the United States and Germany

7.1. Data

We illustrate our methodological contributions with a comparison of relative bipolarisation between the United States and Germany across time. We used data from two long-term income surveys: Panel Study of Income Dynamics (PSID) for the United States and Socio-Economic Panel (SOEP) for Germany. Both surveys are longitudinal, but in order to assess bipolarisation will we use them as repeated cross-sections (using the appropriate weights for cross-sectional data).

In order to ensure comparability by using the same income categories, we resorted to the harmonized Cross-National Equivalent File (CNEF). For individual income comparisons we used labour income before transfers (variable I11110 – Individual Labour Earnings; components: labour earnings, asset flows, private transfers, and private pensions). For household income comparisons, we used household pre-government income (variable I11101), which combines all income before taxes and government transfers across all household members. It was calculated as the sum of total income from labour earnings, asset flows, private transfers, and private pensions. For the sake of completeness, at some points we also took into account household income after taxes and government transfers (variable I11102 for Germany and I111113 for US). However, we note that pre-government income can give a better outlook on labour market income bipolarisation, i.e. before the "smoothing" effect of public transfers.

These data sets allow for an assessment of relative bipolarisation in a relatively long period. In the case of the United States, CNEF data are available for 1970-2009. For Germany, the data cover 1984-2012. In order to maintain comparability, some analyses will be limited to the period 1984-2009.

The compared countries differ significantly in terms of welfare regimes. Germany is a country with much more developed welfare state institutions. According to the typology proposed by Esping-Andersen (1990), it is a conservative (corporatist) welfare state, which is characterized, among other things, by greater extent of income redistribution. By contrast, The United States limits the areas of social policy intervention, following the liberal welfare state model (Esping-Andersen 1990). This involves differences not only in the actually used policy instruments aimed at redistributing income, but also in the perception of guarantees provided by the state to its citizens, and the resulting behavioural consequences; particularly pertaining to the pursuit of income-earning activities. This problem will be analysed in more details hereinafter.

7.2. Relative bipolarisation for individuals

We start with an assessment of relative bipolarisation for individuals. As already mentioned, at an individual level, available data concern pre-government (pre-tax and transfer) income[10] (some benefits, primarily those related to income poverty, are addressed to the household, not individuals). Table 1 provides the descriptive statistics.

[10] All values are given at constant prices for 2016, after adjusting for inflation.

Table 1. Descriptive statistics for individuals

Year	USA				Germany			
	No of obs.	Pre-government income (at 2016 prices)			No of obs.	Pre-government income (at 2016 prices)		
		Mean	Median	Gini		Mean	Median	Gini
1970	6428	35943	28920	0.48				
1971	6622	36452	29623	0.48				
1972	7098	36548	28685	0.48				
1973	7190	37903	30533	0.48				
1974	7585	36199	29204	0.47				
1975	7910	36418	28975	0.47				
1976	8068	37149	29504	0.48				
1977	8216	38339	30636	0.48				
1978	8546	38996	30596	0.48				
1979	9109	37676	29758	0.48				
1980	9397	36151	29111	0.47				
1981	9329	36035	28180	0.47				
1982	9334	36447	29820	0.47				
1983	9291	37402	28889	0.48				
1984	9398	37728	29141	0.48	7168	28292	26420	0.37
1985	9713	39917	30367	0.49	6496	28324	26059	0.39
1986	9682	41249	31716	0.48	6362	28517	26831	0.38
1987	9808	41564	31651	0.48	6373	29314	27354	0.38
1988	9890	42004	32438	0.48	6130	30054	28528	0.38
1989	9989	42392	31910	0.48	5944	30691	28757	0.38
1990	12945	41845	32297	0.48	5846	31547	29300	0.38
1991	12925	41648	31684	0.48	5866	31453	29253	0.38
1992	13298	42826	32546	0.48	8631	28455	25527	0.39
1993	13185	43822	33190	0.48	8254	29706	27003	0.38
1994	10242	47609	35690	0.47	8192	30508	28074	0.37
1995	10274	47371	36180	0.47	8329	30554	28417	0.38
1996	10548	47346	36516	0.47	8257	31521	28770	0.38
1997	8855	46472	35134	0.49	7959	31383	28421	0.38
1998					8621	31648	28946	0.39
1999	9483	50021	37499	0.48	8545	31103	28016	0.40
2000					13961	32569	29390	0.40
2001	10320	55264	40293	0.50	12768	31756	28757	0.41
2002					13971	31298	28091	0.41
2003	10660	54243	40090	0.50	12997	32767	29060	0.42
2004					12468	31951	28230	0.42
2005	10964	58731	40611	0.53	11727	31462	27799	0.42
2006					12256	31591	27366	0.43
2007	11251	60562	41308	0.51	11514	31239	27344	0.43
2008					11016	30848	26077	0.43
2009	11549	63814	42997	0.52	11657	31190	27178	0.43
2010					10592	30915	26385	0.43
2011					11699	30720	26438	0.43
2012					11611	30907	26136	0.44

We observe a significant increase, both in mean and median income, for the US during the whole period, even though the growth rates were quite diverse over time (with average and median income falling in some periods). By contrast, in Germany both the average and the median remained virtually unchanged during 1984-2012.

An important feature, which will be relevant for further discussion, is the increase in relative inequality observed in both countries throughout the period, as measured by the Gini coefficient. It should be noted that in every year, Gini inequality was much higher in the US than in Germany. This observation is part of a broader trend, which was described by Piketty (2014). He stresses higher income inequality in the United States compared with Europe (on the European scale, modern Germany is a country with an average income inequality, higher than in Nordic countries but lower than in southern European countries like Italy, Spain or France[11]).

Changes observed for the level of Gini inequality are *partially* reflected in the results regarding relative bipolarisation. In order to measure relative bipolarisation (Figure 2), we used two indices: $B(Y; 0.5,2,0.5)$, and the index P_4^N proposed by Wang and Tsui (2000), for comparison purposes.[12] As mentioned, changing the values of α (=2) and β (=0.5) affects the extent to which within-partition-group income inequality is reflected in assessment of bipolarisation as a clustering effect. The higher the difference between α and β, the lower the values of the relative bipolarisation index, as the clustering component becomes more negative (see section on decomposition). Changing these parameters, however, has a limited practical impact on our empirical illustration.

Figure 2. Relative bipolarisation for individuals – USA and Germany, pre-government income

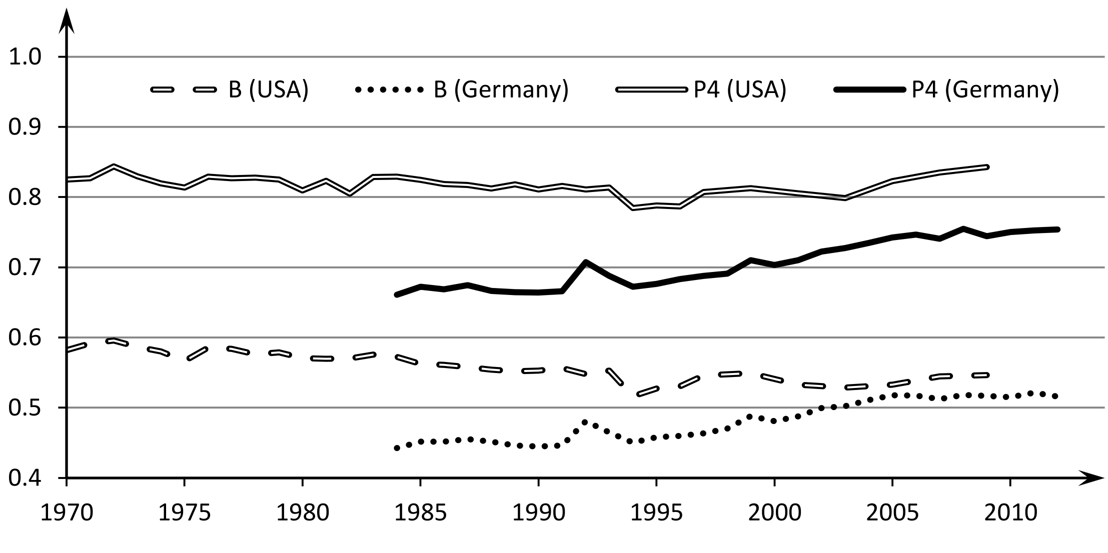

[11] See OECD Database (https://data.oecd.org/inequality/income-inequality.htm).

[12] P_4^N is a median-dependent, rank-independent index which averages the normalized modulus of the distance between each income and the median, where the normalization factor is the median itself. See Wang and Tsui (2000, p. 359).

Figure 2 shows that relative bipolarisation in Germany rose on average between 1984 and 2012. Thus, it approached the level observed in the United States, whose bipolarisation remained relatively constant (index P_4^N), or even decreased slightly (index B).

Figure 2 showed bipolarisation results for $p = 0.5$. However, given the growing inequality in both countries, and the possibility of choosing alternative partition percentiles, we now complement the results with an analysis of the changes in relative bipolarisation at the top of the income distribution. Figure 3 presents measures of relative bipolarisation during 1984-2009 for $p = 0.5$ and $p = 0.95$ (in this latter case, the upper group includes the top 5% income earners). Figure 3's vertical axis measures ratios of B in a given year to the base value of B in 1984.

Figure 3. Trends in relative bipolarisation for USA and Germany, 1984-2009, pre-government income

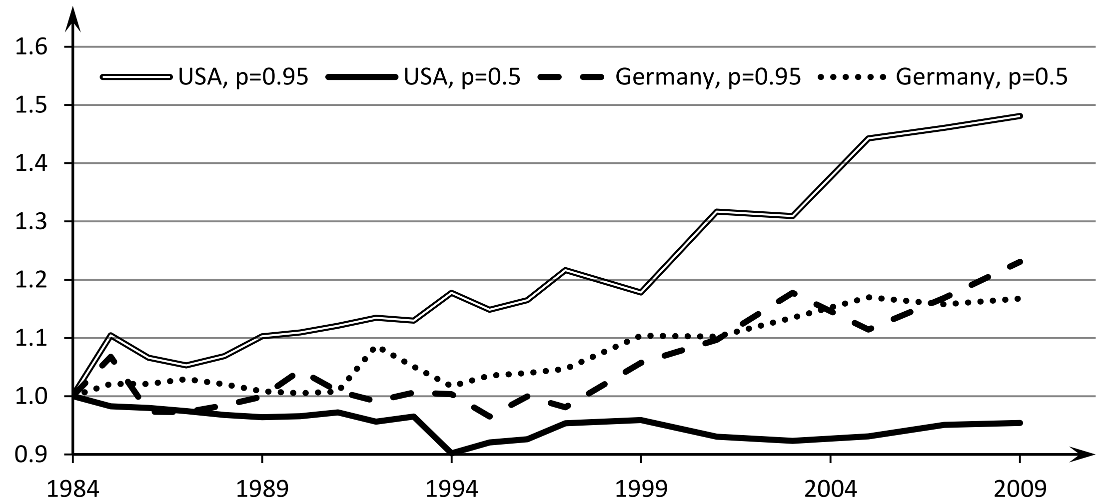

Interestingly, relative bipolarisation increased both in Germany and in the United States for $p = 0.95$, but the increase was faster for the United States. This faster growth is not, however, reflected in the increase in inequality (see Table 1). We can explain this situation by looking at the other series shown in Figure 3. The relative bipolarisation trends for Germany are similar for $p = 0.5$ and $p = 0.95$. This means that changes in relative bipolarisation are similar when either the top group comprises half of all individuals or just the top 5%. This indicates that the rate of income growth increases with income level (and is the highest for the wealthiest people). Both the middle class (people above median income), and the group of the richest individuals, gradually move away from the group of the poor.

In contrast to Germany, whose two bipolarisation series (for $p = 0.5$ and $p = 0.95$) grew at a similar rate, the US series for $p = 0.5$ decreased (while the series for $p = 0.95$ is rapidly increasing). Steadily growing differences between the 5% top earners and the rest of the population were not accompanied by an increase in the difference between the top and bottom halves of income earners (these differences effectively decreased). An explanation can be found in the relative deterioration of the situation of individuals between the 50th and

95th percentile of the income distribution. For a more detailed analysis of this phenomenon we resort to the hybrid Lorenz curves, presented in Figure 4.

Figure 4. Hybrid Lorenz curves for individuals for p=0.5, pre-government income

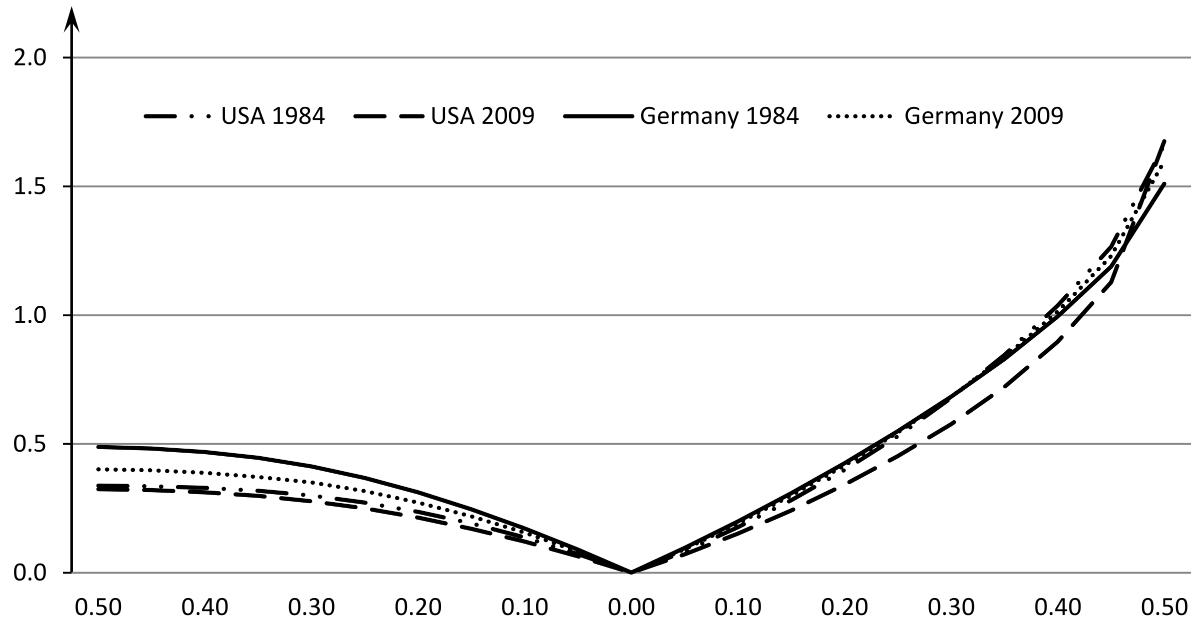

The intersections between all the L-curves mean that there is no dominance relation between any two curves displayed in Figure 4. Therefore, results of the comparison of bipolarisation in US and Germany for individuals between 1984 and 2009 are not fully robust and may depend on the choice of bipolarisation index.

Nevertheless, we can point out some observations from the curves. Firstly, both in Germany and in the US between 1984 and 2009, the relative situation of the poorest half (vis-a-vis the wealthiest half) unequivocally deteriorated, while inequality within this same group decreased.

Secondly, the relative situation of individuals belonging to the highest percentiles has significantly improved in both countries. In the case of the US this improvement concerns, however, only a small group of a few percent of people. Hence we observed in Figure 3 a slight decrease in bipolarisation, when $p = 0.5$ and a rapid increase when $p = 0.95$.

Figure 5. Spread and clustering effects for individuals, pre-government income

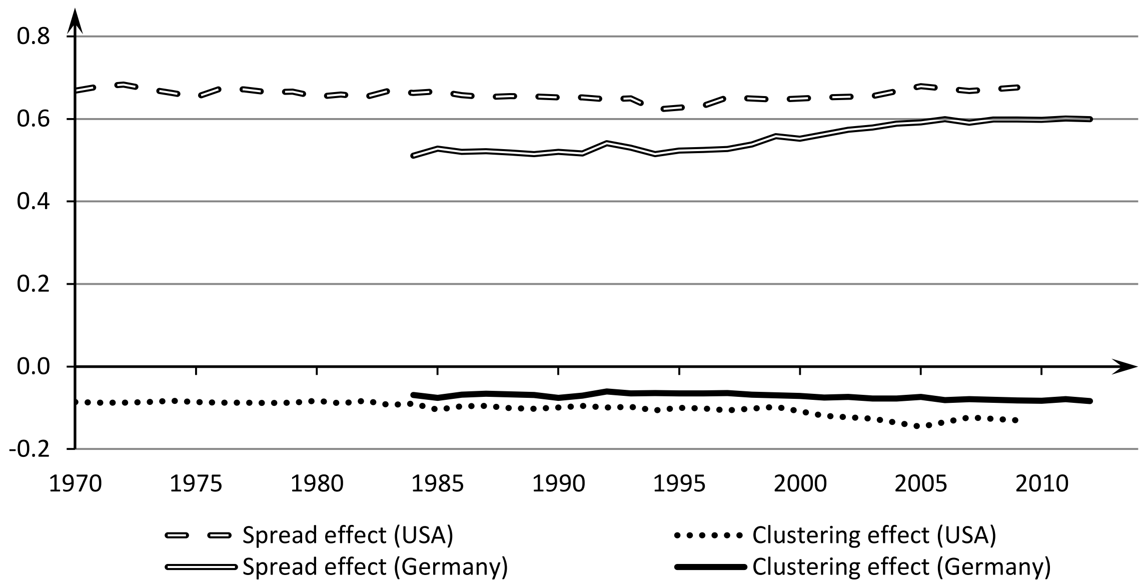

Figure 5 shows the decomposition of B indices for $p = 0.5$ (with $\alpha = 2$ and $= 0.5$). The US shows higher values of both spread and clustering effect (in terms of absolute value) vis-a-vis Germany; i.e. greater values for both the mean distance between the two groups and inequality within each of them, in the US. The combined effect of these trends, however, is a relatively constant level of bipolarisation (see Figure 2; in conjunction with a fixed level of spread effect for the US, it confirms previous considerations on increase in inequality in the upper part of the income distribution). Significantly faster growth of spread effect in Germany, together with changes in clustering effect similar to those that took place in the US, caused the increase in relative bipolarisation in this country and the reduction of the bipolarisation differences between the US and Germany (see Figure 2).

7.3. Bipolarisation for households

Just like individuals, Table 2 shows that pre-government household income's growth rates diverged significantly in both Germany and the US. While mean income increased in both countries (however, much more in the US), median income declined in Germany between 1984 and 2012.

In both countries we observe an increase in the level of inequality during their respective accounting periods (including the common period 1984-2009). By contrast to individual income, however, household income was more unequally distributed in Germany.

Table 2. Descriptive statistics for households

Year	USA				Germany			
	No of obs.	Pre-government income (at 2016 prices)			No of obs.	Pre-government income (at 2016 prices)		
		Mean	Median	Gini		Mean	Median	Gini
1970	4641	57394	50053	0.42				
1971	4831	57533	49767	0.43				
1972	5056	58503	49912	0.43				
1973	5281	58540	50259	0.44				

1974	5511	56887	48673	0.44				
1975	5716	56251	47251	0.45				
1976	5859	55688	46363	0.46				
1977	5997	57810	47887	0.46				
1978	6145	58473	48779	0.45				
1979	6369	58457	47149	0.47				
1980	6529	55842	45995	0.46				
1981	6611	56029	43264	0.48				
1982	6734	55931	44730	0.47				
1983	6833	57044	44941	0.48				
1984	6899	57300	46096	0.48	5624	30909	28193	0.53
1985	7014	61641	47501	0.49	5053	31296	27796	0.54
1986	6998	62646	48667	0.49	4831	32322	29223	0.53
1987	7036	63271	48954	0.49	4771	33389	30464	0.52
1988	7086	64210	48656	0.50	4571	34353	30792	0.53
1989	7096	64620	48348	0.50	4445	34756	30837	0.52
1990	7311	62985	46977	0.50	4401	36260	31573	0.53
1991	7351	60899	45766	0.50	4426	36674	31859	0.52
1992	7531	61367	46069	0.51	6326	33411	28747	0.53
1993	7825	64518	48126	0.51	6298	34348	29408	0.53
1994	9260	65191	45749	0.53	6442	34408	29468	0.53
1995	9058	64343	45617	0.52	6605	34033	28791	0.54
1996	8469	64577	46109	0.52	6525	34798	28571	0.54
1997	6503	68517	49335	0.50	6442	34247	27737	0.54
1998					7264	34430	28363	0.55
1999	6985	69040	48988	0.51	7012	35758	29466	0.54
2000					12582	36799	30174	0.54
2001	7386	78253	58483	0.51	11344	35815	28820	0.54
2002					12055	36079	28262	0.56
2003	7806	73484	56635	0.50	11468	35997	27997	0.56
2004					11207	35437	26573	0.57
2005	7971	79464	55974	0.54	10874	34884	25598	0.57
2006					11895	34275	23650	0.59
2007	8270	80194	55432	0.56	11127	33688	24099	0.58
2008					10544	34110	23515	0.59
2009	8649	82176	57382	0.54	11324	34820	25138	0.58
2010					10335	34158	23443	0.58
2011					11718	34558	23684	0.58
2012					11739	34738	24307	0.58

The increase in the Gini inequality of household income was accompanied by an increase in relative bipolarisation, as shown in Figure 6. The trends in bipolarisation were similar for $B(Y;0.5,2,0.5)$ and P_4^N. Interestingly, despite higher relative bipolarisation of individual income for the US, now relative bipolarisation of household income happened to be higher for Germany.

Figure 6. Relative bipolarisation for households – US and Germany, 1970-2012, pre-government income

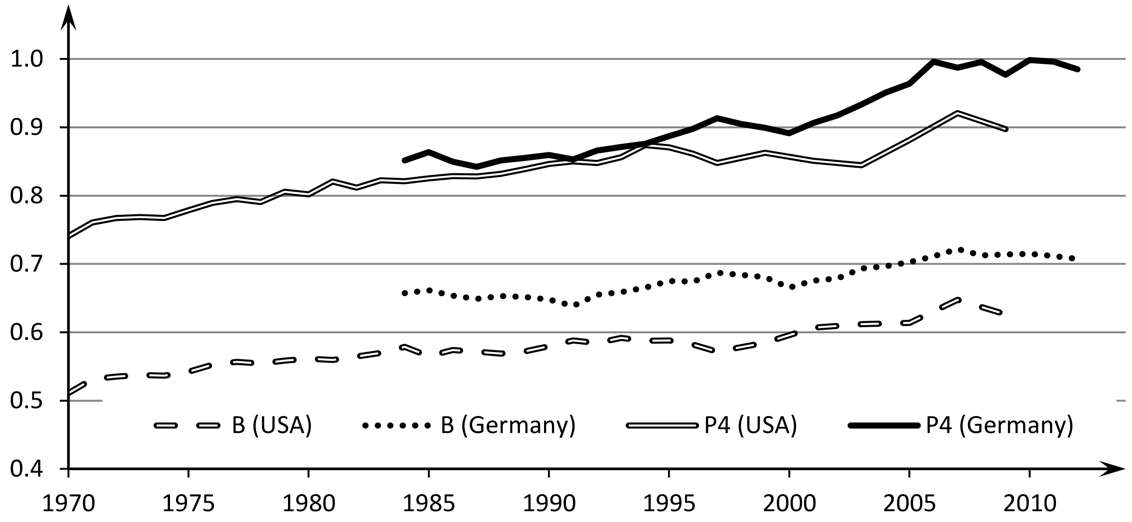

These seemingly contradictory rankings result from the characteristics of the income distribution in both countries. On average, the proportion of households not earning market income is twice as high in Germany compared to the US. It causes a significant increase in bipolarisation, regardless of the shape of income distribution in its upper part. According to Table 3, the ratio of average pre-government income to median income for the highest percentile is systematically higher in the US than in Germany. But the incomes in the lowest deciles are of crucial importance. While average household income is non-zero in the US even in the first decile, in Germany non-zero values occur only from the third decile upward. The reason for this seems to be the prospect of care of state institutions. Significantly smaller support from such institutions in the US (comparing to Germany) suggests that households may have more incentives to seek reliance on market income. And the percentage of households with no income is much lower. The low impact of redistribution on household income in the US is shown by differences between the pre-government and post-government income. Despite the lack of pre-government income in the first and second decile in Germany, post-government relative income (i.e. average income divided by the median) is higher in these groups than in their US' counterparts (despite the latter having higher pre-government income). As a result, in the case of post-government income, comparison of bipolarisation between the US and Germany yields the same results as for individual income: higher relative bipolarisation level in the US than in Germany.

Table 3. Pre- and post-government household income relations, US and Germany

Percentile range	Ratio of average income in a given percentile range to the median income									
	US 1970		US 1984		US 2009		Germany 1984		Germany 2009	
	Pre-gov	Post-gov	Pre-gov	Post-gov	Pre-gov	Post-gov	Pre-gov	Post-gov	Pre-gov	Post-gov
0.00 - 0.10	0.06	0.21	0.01	0.22	0.02	0.17	0.00	0.27	0.00	0.30
0.10 - 0.20	0.26	0.40	0.17	0.40	0.18	0.39	0.00	0.48	0.01	0.49
0.20 - 0.30	0.48	0.58	0.40	0.56	0.42	0.56	0.01	0.64	0.10	0.64
0.30 - 0.40	0.70	0.75	0.61	0.74	0.64	0.71	0.33	0.79	0.31	0.78
0.40 - 0.50	0.90	0.92	0.86	0.91	0.87	0.90	0.86	0.93	0.74	0.92

0.50 - 0.60	1.11	1.09	1.12	1.09	1.15	1.11	1.12	1.09	1.21	1.09
0.60 - 0.70	1.31	1.27	1.39	1.30	1.49	1.36	1.39	1.27	1.63	1.29
0.70 - 0.80	1.59	1.50	1.74	1.55	1.96	1.70	1.70	1.49	2.07	1.54
0.80 - 0.90	1.95	1.81	2.25	1.89	2.65	2.20	2.13	1.79	2.75	1.90
0.90 - 0.95	2.44	2.21	2.91	2.33	3.66	2.92	2.67	2.17	3.62	2.41
0.95 - 0.99	3.22	2.82	3.91	2.97	5.69	4.34	3.48	2.76	5.16	3.29
0.99 - 1.00	5.90	4.41	7.85	5.47	15.41	12.39	6.80	5.01	11.32	6.97

Changes in relative bipolarisation over time are presented in Figure 7. The growth trends for Germany and the US with $p = 0.5$ are very similar, and we observe relatively small increase in relative bipolarisation. Similarly, as in the case of individual income, the largest increase was recorded for the US with $p = 0.95$, whereas Germany's increase at $p = 0.95$ was more moderate. The observed increase in all indices implies rise in the distance between households with the highest income and the rest of the population, which was a phenomenon observed in the context of individual income as well.[13] But in contrast to the situation observed for individual income, relative bipolarisation in household income with $p = 0.5$ also increases in the US. At the household level we do not observe a worsening of the relative situation of people belonging to the middle and upper-middle class (see Table 3).

Figure 7. Trends in bipolarisation for USA and Germany, 1984-2009, pre-government income

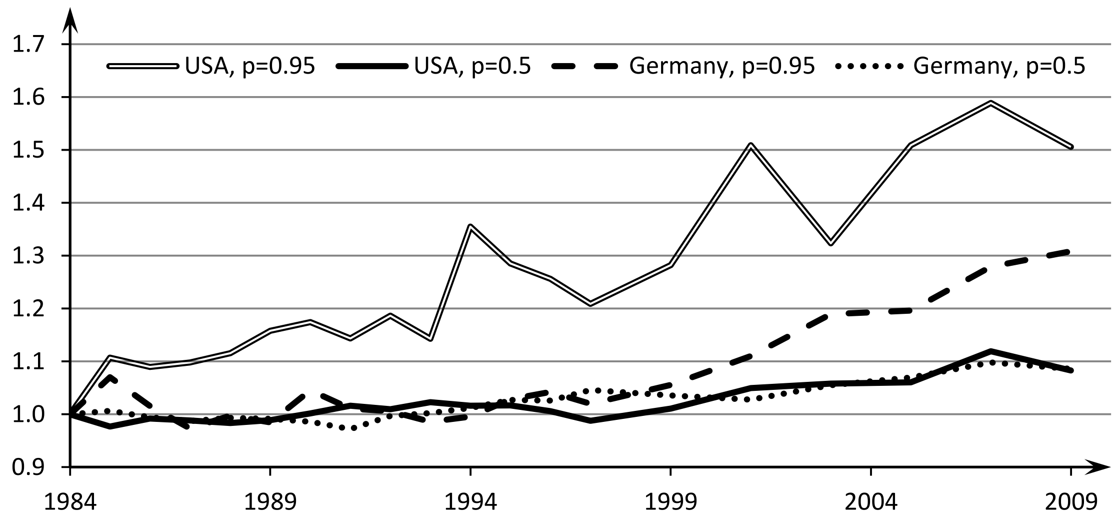

However, due to the intersection of hybrid Lorenz curves (Figure 8), the cross-country comparisons are not independent of the choice of relative bipolarization index.

[13] It is worth stressing that this rise in distances favouring the wealthiest households does not exclusively equate with increase in the spread component. If $p = 0.95$, then between-group distances reflect the spread component. But for $p = 0.5$ distances increase also within the top-half group and it involves a clustering effect besides the spread effect. So, the interpretation in terms of spread and clustering effects depend on the choice of p.

As Figure 8 shows, the RL-curves of both Germany and the US moved downward unambiguously between 1984 and 2009, thereby signalling a worsening of average income among the poorest half relative to the richest. This is consistent with the observed increase in bipolarization in both countries between 1984 and 2009 (Figure 6; although this result depends on the choice of bipolarisation index since the countries' L-curves cross).

The L-curves again show a sharp rise in the highest parts of the income distribution, leading to a strong increase in bipolarisation between the wealthiest households (5%) and the rest of the population.

Figure 8. Relative bipolarisation curves for households with p=0.5, pre-government income

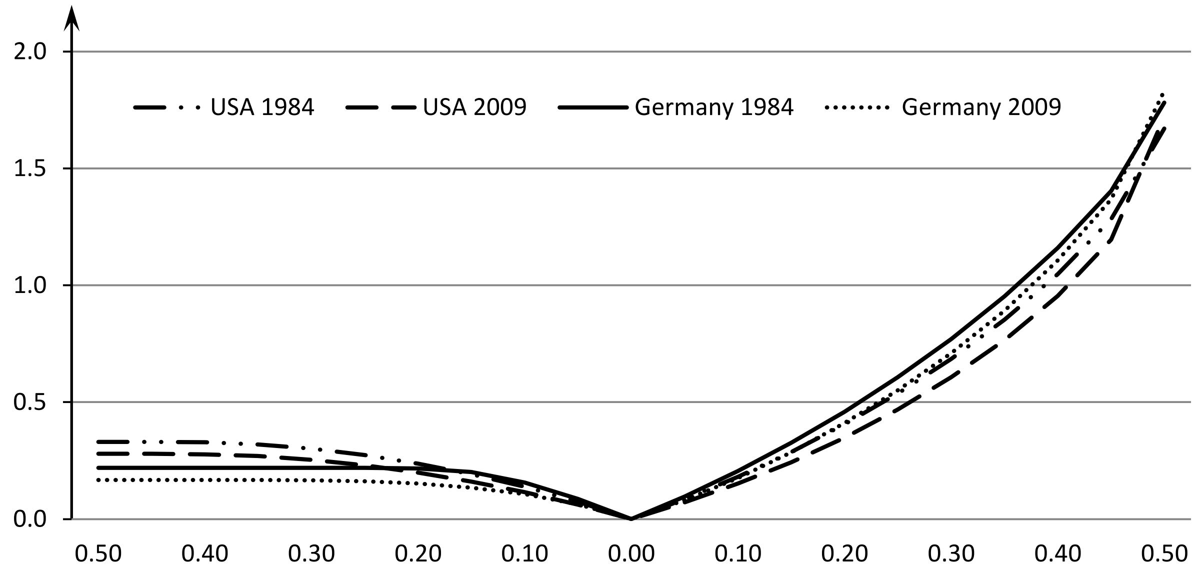

Regarding the RL-curves, it is worth noting that the number of zero income among the worst-off have a direct impact on the length of the horizontal section at the bottom left side of the chart. In the case of Germany zero and very low pre-government income accounted for approximately 30% of the population, hence the very flat course of the entire RL-curve – for both 1984 and 2009. A much smaller percentage of zero and very low income for the US resulted, however, in a much steeper RL-curve.

Figure 9. Spread and clustering effects for households, pre-government income

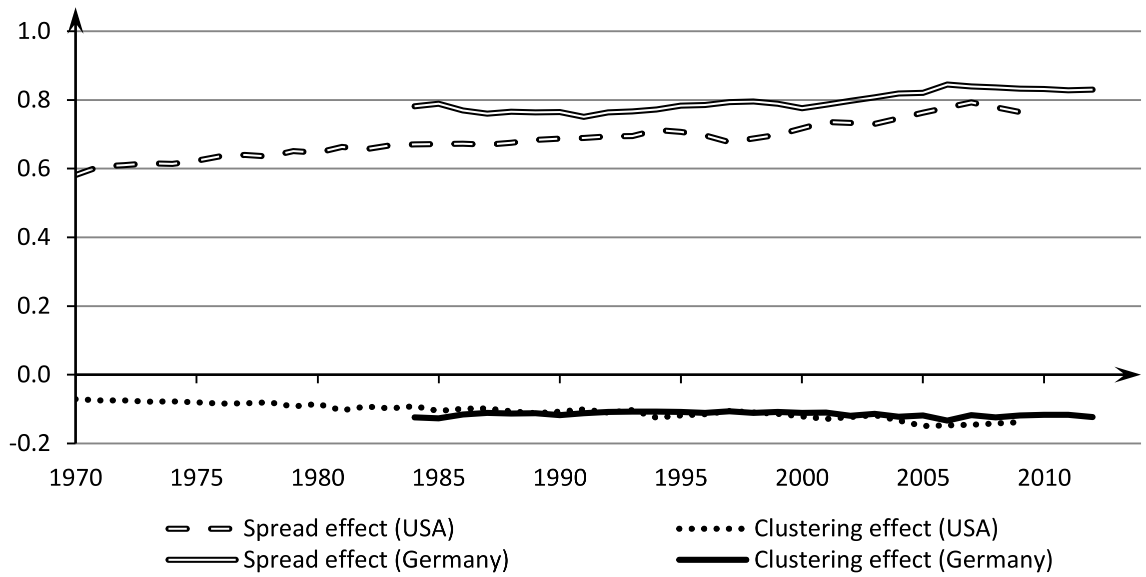

In contrast to the results for the bipolarisation decomposition for individual income (see Figure 5), in the case of household income, the spread effect is higher for Germany than for the US, but the differences slightly decreased in years 1984-2009. Meanwhile, the clustering effect steadily grows in importance in the US (Figure 9), which is associated with the observed increase in inequality (see Table 3). For Germany, the level of clustering effect remained constant through the whole period.

8. Concluding remarks

Our first methodological contribution was the axiomatic characterization of relative bipolarisation indices satisfying the two transfer axioms originally proposed by Foster and Wolfson (2010) together with the defining property of the relative approach: ratio-scale invariance. The characterization drew from the seminal theorem 2 by Bossert and Schworm (2008) and provided a conceptual background for our second contribution, i.e. the introduction of a new class of indices of relative bipolarisation.

As a second methodological contribution, we introduced the first class of indices of relative bipolarisation which are both percentile-independent and partially rank-independent. These indices are based on normalised differences of generalised means and bear the computational advantage of rendering any rank function superfluous (aside from simply requiring the partitioning of the population into two non-overlapping groups). Additionally, we showed that the class of indices is easily decomposable into a spread component and a clustering component. Given the indices' advantages, future research could inquiry into the existence of alternative percentile-independent and rank-independent functional forms.

More specifically, a complete axiomatic characterization of these classes of indices is a desirable pursuit.

For our third contribution, we owe again a debt of intellectual gratitude to the seminal paper by Bossert and Schworm (2008), which developed a median-independent quasi-ordering for absolute bipolarisation measurement. Inspired by their approach, we derived a quasi-ordering for *relative* bipolarisation measurement, a framework which relies on two benchmarks of extreme bipolarisation (i.e. minimum and maximum) unlike others (e.g. absolute). Our quasi-ordering is based on novel *hybrid Lorenz curves* which combine features of both relative and generalised Lorenz curves. Not only are these curves free from the problems caused by reliance on percentiles values in their construction, but they can accommodate any partition of the population into two non-overlapping and exhaustive groups.

Fourthly, we sought to popularise the idea that relative bipolarisation assessments can be performed for any partition of distributions into two groups (i.e. not just identical halves using the median). For that purpose, we introduced the concept of relative *p-bipolarisation*. Relying on the hybrid Lorenz curves, we characterized the very few situations in which one distribution can dominate another one in terms of relative bipolarisation *across the whole percentile domain*.

Our empirical case-study nicely illustrated the relevance and usefulness of our methodological contributions, in addition to being intrinsically interesting. We compared relative bipolarisation in household and individual incomes between the US and Germany across time. Firstly, choosing different group partitions proved relevant in highlighting differences between the two countries. For instance, while individual income bipolarisation grew similarly in Germany for $p = 50\%$ and $p = 95\%$, in the US relative bipolarisation grew very fast with $p = 95\%$, while experiencing a mild *decline* with $p = 50\%$.

Secondly, these two choices of group partitions enabled us to identify and diagnose the relatively unfavourable situation of the upper-middle-class in the United States vis-à-vis the very wealthy and poorer segments of society. Thirdly, the hybrid Lorenz curves revealed that our results were not fully robust to any conceivable choice of relative bipolarisation index. But we should stress that even if we uncovered dominance relationships of relative bipolarisation (we were actually close to these situations of full robustness for the comparison of household income between the two countries in 2009, with $p = 0.5$), we could not generalize these dominance relationships across the whole range of p.

Finally, the hybrid Lorenz curves also proved useful to assess the relative situation of specific groups. For instance, across all income comparisons, the curves allowed us to spot the relative deterioration of the bottom half of the two societies, as well as the significant relative improvement among the wealthiest, between 1984 and 2009.

9. References

Bossert, W. and W. Schworm (2008) "A class of two-group polarization measures" *Journal of Public Economic Theory,* 10(6), 1169–87.

Chakravarty, S. and C. D'Ambrosio (2010) "Polarization orderings of income distributions" *Review of Income and Wealth*, 56(1), 47–64.

Deutsch, J., J. Silber, and M. Hanoka (2007) "On the link between the concepts of kurtosis and bipolarization", *Economics Bulletin*, 4(36), 1–5.

Esping-Andersen, G. (1990) *The Three Worlds of Welfare Capitalism*, Princeton University Press.

Foster, J. and M. Wolfson (2010) "Polarization and the decline of the middle class: Canada and the US", *Journal of Economic Inequality*, 8(2), 247–73.

Lasso de la Vega, C., A. Urrutia, and H. Diez (2010) "Unit consistency and bipolarization of income distributions", *Review of Income and Wealth*, 56(1), 65–83.

Marshall, A., I. Olkin, and B. Arnold (2011) *Inequalities: Theory of Majorization and its applications (Second ed.)*, Springer Series in Statistics. Springer.

Piketty, T. (2014) *Capital in the Twenty-First Century*, Harvard University Press.

Wang, Y.-Q. and K.-Y. Tsui (2000) "Polarization orderings and new classes of polarization indices", *Journal of Public Economic Theory*, 2(3), 349–63.

Wolfson, M. (1994) "When inequalities diverge", *The American Economic Review*, 84(2): 353-8.

Yalonetzky , G. (2014) "Relative bipolarisation Lorenz curve", Economics Bulletin, 34(2): 1103-13.

Yalonetzky, G. (2016a) "The benchmark of maximum relative bipolarisation", ECINEQ WP 2016-419.

Yalonetzky, G. (2016b) "The necessary requirement of median independence for relative bipolarisation measurement", ECINEQ WP 2016-420.

10. Appendix:

Proof of Theorem 1:

Sufficiency: if (a) $B(Y;p)$ is of the form $B(Y;p) = f(\frac{\overline{Y}}{N\mu_Y}, \frac{\underline{Y}}{N\mu_Y})$, (b) $B(Y;p)$ is increasing and Schur-concave in $\overline{Y}$, and (c) $B(Y;p)$ is decreasing and Schur-concave in $\underline{Y}$, then $B(Y;p)$ satisfies axioms (SR), (CI), and (SC).

This can be ascertained by simple inspection. The demonstration provided by Bossert and Schworm (2008, p. 1179) applies, but bearing in mind the transfer-related definition of (SR). As for fulfilment of (SC) we get: $B(\theta Y; p) = f\left(\frac{\theta}{N\theta\mu_Y}\overline{Y}, \frac{\theta}{N\theta\mu_Y}\underline{Y}\right) = f\left(\frac{\overline{Y}}{N\mu_Y}, \frac{\underline{Y}}{N\mu_Y}\right) = B(Y; p)$.

Necessity: If $B(Y; p)$ satisfies axioms (SR), (CI), and (SC), then it must be the case that: (a) $B(Y; p)$ is of the form $B(Y; p) = f(\frac{\overline{Y}}{N\mu_Y}, \frac{\underline{Y}}{N\mu_Y})$, (b) $B(Y; p)$ is increasing and Schur-concave in $\overline{Y}$, and (c) $B(Y; p)$ is decreasing and Schur-concave in $\underline{Y}$.

Parts (b) and (c) are ascertained following the demonstration by Bossert and Schworm (2008, p. 1179). For part (a) we start by restating (SC): $B(\theta Y; p) = f(\theta\overline{Y}, \theta\underline{Y}) = f(\overline{Y}, \underline{Y}) = B(Y; p)$. Now let $s \subseteq Y$. In order to have $f(\theta\overline{Y}, \theta\underline{Y}) = f(\overline{Y}, \underline{Y})$ we clearly need a function $f\left(\frac{\overline{Y}}{g(s)}, \frac{\underline{Y}}{h(s)}\right)$ where $g(s)$ and $h(s)$ must be linearly homogenous (in order to get θ cancelled out).

But why must it also be the case that $g(s) = h(s) = N\mu_Y$, specifically? Because any other alternative linearly homogeneous function, and/or using just a subset of Y, would lead to a violation of either (SR), (CI), or both. This is so because, for any other linearly homogeneous function, and/or using just a subset of Y, one can find at least one transfer that would alter the denominators of $\frac{\overline{Y}}{g(s)}$ and/or $\frac{\underline{Y}}{h(s)}$. Then $B(Y; p)$ could be in violation of (SR), (CI), or both.

■

Proof of Proposition 1:

Satisfaction of SD:

Consider a Pigou-Dalton transfer involving incomes i and j such that $y_i < y(p) < y_j$, where $y(p)$ is the quantile corresponding to percentile p, hence i and j belong in different parts of the distribution. The change in the index due to a transfer of amount γ is:

$$\frac{\partial B}{\partial \gamma} = \frac{1-p}{\mu_Y}\left[-\left[\frac{1}{N(1-p)}\sum_{k=Np+1}^{N} y_k^{\beta}\right]^{\frac{1}{\beta}-1} y_j^{\beta-1} - \left[\frac{1}{Np}\sum_{k=1}^{Np} y_k^{\alpha}\right]^{\frac{1}{\alpha}-1} y_i^{\alpha-1}\right] < 0 \quad (9)$$

Note that $\frac{\partial B}{\partial \gamma} < 0$ is true for any value (α, β) in which both parameters are non-zero.

Satisfaction of CI:

Here is where the restriction $\alpha > 1 > \beta$ plays a prominent role. Consider a Pigou-Dalton transfer involving incomes i and j such that $y_i < y_j < y(p)$, where $y(p)$ is the quantile corresponding to percentile p, hence i and j belong in the same bottom part of the distribution. The change in the index due to a transfer of amount γ is:

$$\frac{\partial B}{\partial \gamma} = \frac{1-p}{\mu_Y}\left[\frac{1}{Np}\sum_{k=1}^{Np} y_k^{\alpha}\right]^{\frac{1}{\alpha}-1} \left[y_j^{\alpha-1} - y_i^{\alpha-1}\right] \quad (10)$$

Given that $y_i < y_j$, it is clear that $\frac{\partial B}{\partial \gamma} > 0$ if and only if $\alpha > 1$.

Now consider a Pigou-Dalton transfer involving incomes i and j such that $y(p) < y_i < y_j$, where $y(p)$ is the quantile corresponding to percentile p, hence i and j belong in the same bottom part of the distribution. The change in the index due to a transfer of amount γ is now:

$$\frac{\partial B}{\partial \gamma} = \frac{1-p}{\mu_Y}\left[\frac{1}{N(1-p)}\sum_{k=Np+1}^{N} y_k^{\beta}\right]^{\frac{1}{\beta}-1}\left[y_i^{\beta-1} - y_j^{\beta-1}\right] \quad (11)$$

Now, given that $y_i < y_j$, it is clear that $\frac{\partial B}{\partial \gamma} > 0$ if and only if $\beta < 1$.

Satisfaction of N, part (a):

If $Y \in \mathcal{E}$, then it will be the case that: $\overline{\mu}(Y;p,\beta) = \underline{\mu}(Y;p,\alpha)$. Hence $B(Y;p,\alpha,\beta) = 0$. Conversely, if $B(Y;p,\alpha,\beta) = 0$ then it must be the case that: $\overline{\mu}(Y;p,\beta) = \underline{\mu}(Y;p,\alpha)$. The latter is only possible $Y \in \mathcal{E}$. In fact, this is true for the whole subset of non-null (α,β), i.e. not just for $\alpha > 1 > \beta$.

Satisfaction of N, part (b):

If $Y \in \mathcal{B}_p$, then it will be the case that $\underline{\mu}(Y;p,\alpha) = 0$ and $\overline{\mu}(Y;p,\beta) = y > 0$. Likewise, $\mu_Y = y(1-p)$. Therefore $B(Y;p,\alpha,\beta) = 1$ if $Y \in \mathcal{B}_p$. This will be true for the whole subset of non-null (α,β). Meanwhile, if $B(Y;p,\alpha,\beta) = 1$ then, from $B(Y;p,\alpha,\beta) \equiv \frac{1-p}{\mu_Y}[\overline{\mu}(Y;p,\beta) - \underline{\mu}(Y;p,\alpha)]$ (3), we can deduce that $\overline{\mu}(Y;p,1) + \underline{\mu}(Y;p,1) = 2(1-p)[\overline{\mu}(Y;p,\beta) - \underline{\mu}(Y;p,\alpha)]$. Rearranging we get: $\overline{\mu}(Y;p,1) - 2(1-p)\overline{\mu}(Y;p,\beta) = -\underline{\mu}(Y;p,1) - 2(1-p)\underline{\mu}(Y;p,\alpha)$. Now this last expression could be satisfied with different distributions, depending on the values of (α,β). However, if $\alpha > 1 > \beta$, then $\overline{\mu}(Y;p,1) - 2(1-p)\overline{\mu}(Y;p,\beta) > -\underline{\mu}(Y;p,1) - 2(1-p)\underline{\mu}(Y;p,\alpha)$, unless $Y \in \mathcal{B}_p$, in which case: $\overline{\mu}(Y;p,1) - 2(1-p)\overline{\mu}(Y;p,\beta) = -\underline{\mu}(Y;p,1) - 2(1-p)\underline{\mu}(Y;p,\alpha)$. Therefore, if $\alpha > 1 > \beta$, then $B(Y;p,\alpha,\beta) = 1$ implies that $Y \in \mathcal{B}_p$.

Conversely, if $Y \notin \mathcal{B}_p$, then we could have $B(Y;p,\alpha,\beta) < 1$ if, for instance, the top part is egalitarian, and the bottom part has only one positive income. Likewise, we could get $B(Y;p,\alpha,\beta) > 1$, if the bottom part has only zero incomes, the top part is not egalitarian, and $\beta > 1$. Hence $Y \in \mathcal{B}_p$ is also necessary for $B(Y;p,\alpha,\beta) = 1$. ■

Proof of Theorem 2:

The role of SY is trivial to prove. Now since we want the bipolarisation measures and hybrid Lorenz curves to satisfy PP then we should work with continuous versions of L and R in order to compare the quantiles of distributions with different population sizes:

$$L(Y, p, 1; k) = \frac{1}{1-p} \int_p^k \frac{y(q)}{\mu_Y} dq \quad \forall k \in [p, 1] \quad (12)$$

$$R(Y, 0, p; k) = \frac{1}{p} \int_0^k \frac{y(p-q)}{\mu_Y} dq \quad \forall k \in [0, p] \quad (13)$$

Then rephrase the theorem the following way:

Theorem 2a: $B(X; p) > B(Y; p)$ for all B satisfying SY, PP, SC, SD, CI and N if and only if (i) $L(X, p, 1; k) \geq L(Y, p, 1; k)\ \forall k \in [p, 1]$, with at least one strict inequality; and (ii) $R(X, 0, p; k) \leq R(Y, 0, p; k)\ \forall k \in [0, p]$, with at least one strict inequality.

Now in order to prove the theorem it will be convenient to prove the following two propositions:

Proposition 2: $B(X; p) > B(Y; p)$ for all B satisfying SY, PP, SC, SD/SR, CI/CR and N if and only if X can be obtained from Y through a sequence of operations involving: (i) multiplications by scalars; (ii) regressive transfers across percentile p; and (iii) Pigou-Dalton transfers on either side of the percentile p.

Proof:

First we prove that obtaining X from Y through the specified sequence leads to $B(X; p) > B(Y; p)$ for any index satisfying SY, PP, SC, SD/SR, CI/CR and N. Let $Z = \lambda Y$ where λ is a positive scalar. Then we know that for any index satisfying SC we have: $B(Y; p) = B(Z; p)$. Now we obtain distribution W from Z using regressive transfers across percentile p. Then we know that for any index satisfying SR we have: $B(W; p) > B(Z; p)$. Finally, we obtain X from W using Pigou-Dalton transfers on either side of p. Then we know that for any index satisfying CR we have: $B(X; p) > B(W; p)$. Hence it is clear, that through this sequence we get $B(X; p) > B(Y; p)$ for any index satisfying SC, SD/SR and CI/CR. Meanwhile if $Y \in \mathcal{E}$, then any other distribution can be transformed into Y through a sequence of only scalar multiplications and Pigou-Dalton transfers across the dividing percentile. Therefore $B(X; p) > B(Y; p)$ for any index satisfying the first part of N. As for the second part, note that any $X \in \mathcal{B}_p$ can be obtained from Y through an appropriate sequence of scalar multiplication, regressive transfers (basically rendering all the incomes in the bottom part equal to 0), and finally Pigou-Dalton transfers among incomes in the top part. Therefore $B(X; p) > B(Y; p)$ for any index satisfying the second part of N.

The second step requires proving that if X is obtained from Y through an alternative sequence then it will not be the case that $B(X; p) > B(Y; p)$ for all B satisfying SY, PP, SC, SD/SR, CI/CR and N. For example, take any Y, then perform one Pigou-Dalton transfer across percentile p, followed by one Pigou-Dalton transfer on the bottom part and another Pigou-Dalton transfer on the top part. Then it should not be difficult to find two bipolarisation indices, B^1 and B^2, such that: $B^1(X; p) > B^1(Y; p)$ and $B^2(X; p) < B^2(Y; p)$.[14] ■

[14] Numerical examples are available upon request.

Proposition 3: (i) $L(X, p, 1; k) \geq L(Y, p, 1; k)\ \ \forall k \in [p, 1]$, with at least one strict inequality; and (ii) $R(X, 0, p; k) \leq R(Y, 0, p; k)\ \ \forall k \in [0, p]$, with at least one strict inequality, if and only if X can be obtained from Y through a sequence of operations involving: (i) multiplications by scalars; (ii) regressive transfers across percentile p; and (iii) Pigou-Dalton transfers on either side of the percentile p.

Proof:

First we prove that if X can be obtained from Y through the above sequence of operations then it must be the case that: (i) $L(X, p, 1; k) \geq L(Y, p, 1; k)\ \ \forall k \in [p, 1]$, with at least one strict inequality; and (ii) $R(X, 0, p; k) \leq R(Y, 0, p; k)\ \ \forall k \in [0, p]$, with at least one strict inequality. This can be done, by evaluating one operation at a time. Then we can compound the effects (as they all work in similar directions):

(i) Scalar multiplication: Since the hybrid Lorenz curves are scale invariance then if X is obtained from Y through a scalar multiplication then we get: $L(X, p, 1; k) = L(Y, p, 1; k)\ \ \forall k$ and $R(X, 0, p; k) = R(Y, 0, p; k)\ \ \forall k$.

(ii) Regressive transfers across percentile p: Imagine the transfer involves $y(m)$ and $y(r)$, such that: $m < p < r$. Then we will have: $L(X, p, 1; k) = L(Y, p, 1; k)\ \forall k \in [p, r]$ and $L(X, p, 1; k) > L(Y, p, 1; k)\ \forall k \in [r, 1]$. Meanwhile we will have: $R(X, 0, p; k) = R(Y, 0, p; k)\ \ \forall k \in [0, m]$ and $R(X, 0, p; k) < R(Y, 0, p; k)\ \forall k \in [m, p]$.

(iii) Pigou-Dalton transfers on either side of the percentile p: Imagine the transfer involves $y(r)$ and $y(s)$, such that: $p < r < s$. Then we will have: $L(X, p, 1; k) = L(Y, p, 1; k)\ \forall k \in [p, r] \cup [s, 1]$ and $L(X, p, 1; k) > L(Y, p, 1; k)\ \forall k \in [r, s]$. A similar situation will occur if the transfer involves $y(m)$ and $y(n)$, such that: $m < n < p$.

The second step requires proving that if i) $L(X, p, 1; k) \geq L(Y, p, 1; k)\ \ \forall k \in [p, 1]$, with at least one strict inequality; and (ii) $R(X, 0, p; k) \leq R(Y, 0, p; k)\ \ \forall k \in [0, p]$, with at least one strict inequality, then X can be obtained from Y through the above sequence of operations. Essentially, we have to prove that the hybrid Lorenz curve of X can be obtained from that of Y using the sequence, which should be designed like this: First, multiply Y by $\frac{\mu_X}{\mu_Y}$ and obtain distribution $Z = \frac{\mu_X}{\mu_Y} Y$. Now Z has the mean of X, but the same hybrid Lorenz curves as Y, due to the fulfilment of scale invariance. Since we know that $L(X, p, 1; 1) \geq L(Z, p, 1; 1)$ and $R(X, 0, p; p) \leq R(Z, 0, p; p)$ then it must be the case that $\underline{\mu}(Z; p, 1) > \underline{\mu}(X; p, 1)$ and $\overline{\mu}(Z; p, 1) < \overline{\mu}(X; p, 1)$. Hence the next step is to perform a sequence of regressive transfers from the bottom part of Z to its top part. More specifically, the aim is to produce a new distribution W from Z such that: $L(X, p, 1; 1) = L(W, p, 1; 1)$ and $R(X, 0, p; p) = R(W, 0, p; p)$, i.e. the ends of the two curves of W coincide with the respective ends of X. The most natural starting point is a regressive transfer from the lowest income to the highest income. The sequence and amount of each regressive transfer needs to meet some restrictions which produce an algorithm. For instance, after every regressive transfer it needs to be the case that $L(X,p,1;k) \geq L(W,p,1;k)$ and $R(X,0,p;k) \leq R(W,0,p;k)$. Likewise, if necessary, the second lowest and/or the second highest income may need to be involved and so forth. Finally, depending on the situation it could be the case that, after the regressive transfers, the ending segments (not just the ending points) of the pairs of curves coincide, i.e. $L(X,p,1;k) = L(W,p,1;k)\ \forall k \in [\bar{k},1]$ with $\bar{k} > p$ and $R(X,0,p;k) = R(W,0,p;k)\ \forall k \in [\bar{k},p]$ with $\bar{k} > 0$. What matters is that the ending points, or ending segments, of the pairs of curves overlap. Once we reach this stage then we finally obtain X from W using Pigou-Dalton transfers involving percentiles in the interval $[p,\bar{k}]$ for the top part and percentiles in the interval $[0,\bar{k}]$ for the bottom part. We know that we can obtain X from W at this stage with these transfers due to a continuous version of Muirhead's theorem (Marshall et al., 2011, p.7-8). ■

By now, it should be apparent that by proving Proposition 2 and Proposition 3, thereby establishing an equivalence relationship between the index condition, the hybrid Lorenz curve condition and the sequential derivation condition, we have essentially proven Theorem 2. ■

Proof of Theorem 3:

In the proof of Theorem 2, we established that dominance of A over B, based on the p percentile implies that A can be obtained from B through a sequence of operations involving: (i) multiplications by scalars; (ii) regressive transfers across percentile p; and (iii) Pigou-Dalton transfers on either side of the percentile p. Now consider the case of $A, B \notin \mathcal{E}$. If A dominates B then we will need both regressive transfers (unless $(A,p,1;1) = L(B,p,1;1)$) and Pigou-Dalton transfers on at least one side of p, in order to obtain A from B. Now imagine that the sequence to obtain A from B required a Pigou-Dalton transfer involving percentiles m and n, such that $m < n < p$. This was then a Pigou-Dalton transfer in the bottom part of the distribution. However, if now we choose a different dividing percentile, γ, such that $m < \gamma < n$, then the Pigou-Dalton transfer takes place across the partition, i.e. between the two parts. Instead of producing a clustering effect, as was the case when p was dividing the distribution, now this transfer decreases the spread between the two parts. Instead of contributing to produce a more bipolarised distribution (A) it is offsetting any previous sequential transfers which were increasing bipolarisation. Moreover, it could well happen that with γ now the sequence of transfers is inconsistent, in the sense that it may include clustering-increasing with spread-decreasing transfers, therefore generating curve-crossing! Whichever the case, dominance is no longer maintained when γ replaces p.

The above proof will hold true for most comparisons, but there are two main exceptions wherein actually dominance over the whole percentile domain is possible:

(i) $A \notin \mathcal{E} \land B \in \mathcal{E}$: When one of the pairs of distributions is egalitarian then clearly the other non-egalitarian distribution dominates for every relevant dividing percentile (i.e. $p \in]\frac{2}{N}, 1-\frac{2}{N}[$). The reason is that once we choose a percentile, we can obtain B from A through a sequence of scalar multiplications and just Pigou-Dalton transfers across the chosen percentile (some incomes may switch between parts in the process). Change the percentile, and the sequence can be implemented again, from A to B.

(ii) A was obtained from B through a sequence of regressive transfers involving, each time, one income in percentile $q \leq \frac{2}{N}$ and one income in percentile $r \geq 1-\frac{2}{N}$: In this narrow case the dominance of A holds over the whole relevant percentile domain because the regressive transfer is always spread-increasing, i.e. never cluster-decreasing. ■