Inequalities and Identities

Arjun Jayadev and Sanjay G. Reddy

Abstract We introduce concepts and measures relating to inequality between identity groups. We define and discuss the concepts of *Representational Inequality*, *Sequence Inequality,* and *Group Inequality Comparison*. *Representational Inequality* captures the extent to which an attribute is shared between members of distinct groups. *Sequence Inequality* captures the extent to which groups are ordered hierarchically. *Group Inequality Comparison* captures the extent of differences between groups. The concepts can be used to interpret segregation, clustering, and polarization in societies.

> Civil paths to peace also demand the removal of gross economic inequalities, social humiliations and political disenfranchisement, which can contribute to generating confrontation and hostility. Purely economic measures of inequality do not bring out the social dimension of the inequality involved. For example, when the people in the bottom groups in terms of income have different non-economic characteristics, in terms of race (such as being black rather than white), or immigration status (such as being recent arrivals rather than older residents), then the significance of the economic inequality is substantially magnified by its "coupling" with other divisions, linked with non-economic identity groups.
>
> Amartya Sen, *The Guardian*, Friday November 9, 2007

> A highly ethnically diversified society may generate tensions in the society which ultimately may lead to conflicts…. [An] objective of the society should, therefore, be to make ethnic diversity (hence ethnic polarization) as low as possible.
>
> Satya Chakravarty and Bhargav Maharaj, 2009

This is a modified version of various versions of a working paper by the same title uploaded on the ISERP website, https://academiccommons.columbia.edu/doi/10.7916/D86979DR and on the SSRN website: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=1162275. We are the copyright owners on these.

A. Jayadev (✉)
Azim Premji University, Bengaluru, India
e-mail: arjun.jayadev@apu.edu.in

S. G. Reddy
New School for Social Research, New York, USA
e-mail: reddysanjayg@gmail.com

I. Dasgupta and M. Mitra (eds.), *Deprivation, Inequality and Polarization*,
Economic Studies in Inequality, Social Exclusion and Well-Being,
https://doi.org/10.1007/978-981-13-7944-4_7

1 Introduction

For Satya Chakravarty, whose contributions to scholarship and justice we honor here, the salience of group-based membership for evaluating social outcomes is clear. His remarks, along with those of Amartya Sen above, suggest that the interpersonal differences in advantages can be of greater concern when associated systematically with membership of groups. This can be true for two reasons. First, the avoidance of systematically arising intergroup differences may be of intrinsic importance from the perspectives of justice and fairness in the distribution of goods and opportunities. Second, the fact that there exist distinct groups in society and that these groups exhibit intergroup differences may have instrumental significance from the standpoint of their impact on social goods such as peace, stability, or economic growth

The concern with the intrinsic significance of intergroup differences has centered on the degree to which "morally irrelevant" characteristics of a person (such as belonging to a given race, sex, caste, or other groups as a result of birth) should be permitted to determine her or his life chances.[1] Such a motivation is distinct from one based on the idea that social goods or "bads" may be generated by intergroup differences in economic and social achievements, and that intergroup differences may be relevant for that reason. A body of literature in economics and other social sciences has explored this instrumental concern.[2] Both concerns have led to the development of a growing literature that has identified and empirically examined such concepts as "horizontal inequality", segregation, polarization, and related ideas about differences between groups

That a multitude of concepts concerning intergroup difference has been proposed is not entirely surprising because such differences can be understood as arising in more than one way. For example, studies on segregation focus on the degree to which members of different groups share a location, occupation, or other attribute while studies on horizontal inequality focus on the extent of difference in the income or other achievements of separate groups. In both cases, however, the subject of interest is the degree of unevenness or inequality in the possession of attributes between groups. The goal of our paper is to elucidate some distinct ways in which intergroup differences can be conceived, which encompass but are not restricted to the concerns of these existing approaches.

A common underlying concern in analyses of intergroup differences is the degree to which distinct groups are systematically over- or under-represented in their possession of various attributes (levels of income or health, club membership, political office, etc.). In this paper, we introduce the concept of *Representational Inequality*

---

[1]For a review of these debates, see, e.g., Roemer (1996), and Sen (1992). Arguments that societies should be organized so as to limit the consequences of being born into a particular position as such include those of "luck egalitarians" such as Arneson (1989), Cohen (1989), Dworkin (2000), Rawls (1971) and Roemer (1996). Egalitarians of other kinds (e.g., those concerned with relational equality, such as Anderson 1999) may come to similar conclusions for different reasons.

[2]For some examples, see, e.g., Stewart (2001), Alesina et al (2003), Alesina and La Ferrara (2000, 2002), Montalvo and Reynal Querol (2005), Miguel and Gugerty (2005), and Østby (2008).

(*RI*) as a way to capture this concern. This concept describes the extent to which a given attribute (for instance, a level of income or health, or right or left-handedness) is shared by members of distinct groups. It can be used to measure the degree of "segregation" of distinct identity groups in the attribute space.[3]

When individuals can be ordinally ranked in relation to an attribute (such as income or health but not right- or left-handedness), we may be interested not only in how segregated or separated each identity group is in terms of their achievements, but in some measure of their relative positions in the ranking. *Sequence Inequality (SI)*, understood as the degree to which members of one group are placed higher in a given hierarchy than those from another, captures this concern. Such a concept provides an intuitive framework for understanding the degree of "clustering" of various identity groups in distinct sections of a hierarchy.[4]

When individuals' level of achievement can also be cardinally identified for an attribute (as for income but not for right- or left-handedness), the *distance* between groups' attribute levels may be of interest. We may identify a distinction between two different concepts, which we term, respectively, *Group Inequality Comparison (I)* and *Group Inequality Comparison (II)* and abbreviate as *GIC (I)* and *GIC (II).* The concept of *Group Inequality Comparison (I)* involves a comparison of counter-factuals. Specifically, it is derived by comparing the inequality arising in a society in which all of the members of a group are assigned a representative income for that group and the total interpersonal inequality in a society. This concept is concerned with identifying the extent to which between-group inequality "accounts for" overall inequality in society. *Group Inequality Comparison (II)* by contrast measures only the inequality arising in the first situation, i.e., that in a society in which all the members of a group are assigned a representative income for that group. This latter concept is concerned with the absolute magnitude of the inequality generated by between-group inequality.

Our purpose in this paper is twofold. We seek not only to clarify the concepts described above and thus to recognize the complexities of intergroup differences but also to show that combining these concepts can be helpful in characterizing intergroup differences taken as a whole. "Polarization",[5] understood to involve the collection of like elements and the separation of such collections of like elements from one another, can be fruitfully described as involving the simultaneous presence of between-group differences of different kinds. The combination of Representational Inequality with Sequence Inequality alone provides a measure of what might be termed "Ordinal Polarization." Combining Group Inequality Comparison (of either type I or type II)

---

[3] Segregation is defined by the Oxford English Dictionary, *inter alia,* as "The separation of a portion of portions of a collective or complex unity from the rest; the isolation of particular constituents of a compound or mixture."

[4] A cluster is defined by the Oxford English Dictionary, *inter alia*, as "A collection of things of the same kind…growing closely together; a bunch… a number of persons, animals, or things gathered or situated close together; an assemblage, group, swarm, crowd."

[5] The Oxford English Dictionary defines the verb "polarize" as "To accentuate a division within (a group, system, etc.); to separate into two (or occas. several) opposing groups, extremes of opinion, etc."

with these other two indices can provide a richer index of Polarization applicable to the case in which the attribute is cardinally measurable as well. *Our purpose is not to provide a unique characterization of a single measure of polarization, but rather to show that a broad class of measures of polarization can be derived from a simple set of unexceptionable axioms concerning different types of between-group differences and their combination.*

The concept of polarization that we employ here is distinct from that developed in the preponderance of the existing literature in that it draws on information about the identity groups to which those who possess distinct attributes belong. In contrast, the existing frameworks generally employ a "collapsed" framework in which the level of the attribute (typically income) defines the identity group (Esteban and Ray 1994; Duclos et al. 2004). In these frameworks, polarization of an income distribution is understood to involve "identification" between individuals possessing a certain level of income and "alienation" between those individuals and others possessing different incomes. In our framework, in contrast, polarization of an income distribution is understood to involve segregation of individuals belonging to distinct identity groups at certain levels of income and the separation of these groupings of individuals in the income space from other groupings of individuals possessing distinct identities.

2 Part I: Concepts of Group Inequality

One approach to evaluating intergroup differences is to construct a measure of overall group advantage or disadvantage for each group prior to assessing the differences in these overall measures.[6] Although there can be advantages to such an approach, it can obscure the diverse aspects of intergroup difference (by reducing intergroup differences to inequalities in a single dimension). We accordingly explicitly identify here three distinct concepts of intergroup difference, and a fourth which builds upon them.

2.1 Representational Inequality

We define a situation of representational inequality as occurring when, for some attribute and some identity group, the proportion of the group possessing the attribute is either greater or less than the proportion of the group in the overall population. To provide some graphical intuition for this idea, consider the distribution of income among different groups in a society that consists of 50 percent whites and 50 percent blacks. Figure 1 depicts the situation in which there is no representational inequality. The location of each bar on the horizontal axis represents an income level ordered from lowest to highest and the proportion of persons possessing that income of either

[6]See Jayaraj and Subramanian (2006) for an example of such an approach.

group is represented through shading. At all levels of income, blacks and whites are represented in equal proportion to their share of the population as a whole (i.e., one half each). Any deviation from such equiproportionality leads to a situation of representational inequality. Such a situation is depicted in Fig. 2, in which at certain levels of income blacks or whites comprise a larger or smaller proportion of the individuals possessing that level of income than they do in the population.

While the situation depicted in Fig. 2 is one of the representational inequalities, both groups are represented at all the incomes. In contrast, Fig. 3 depicts a situation in which at each level of income there is *complete* segregation, in the sense that at each level of income there is one and only one identity group represented. It may be noted that although this is a situation of complete segregation, the incomes at which whites and blacks appear are evenly interspersed. We depict this example to make sharp the distinction between segregation and clustering as we use the terms. The former refers to a situation in which those possessing a specific attribute (in this case an income level) belong disproportionately to a group. The latter refers to a situation in which the attributes disproportionately possessed by members of a group are sited together in a certain part of an attribute hierarchy (in this case the income spectrum).

The concept of representational inequality clearly need not be restricted to a scenario in which the attribute is cardinally orderable. Thus, for example, we can apply the principle in an equally straightforward manner to unordered attributes such as location of residence, or membership in distinct clubs or legislatures. If instead of income brackets, each bar referred to a distinct legislature in a federal country,

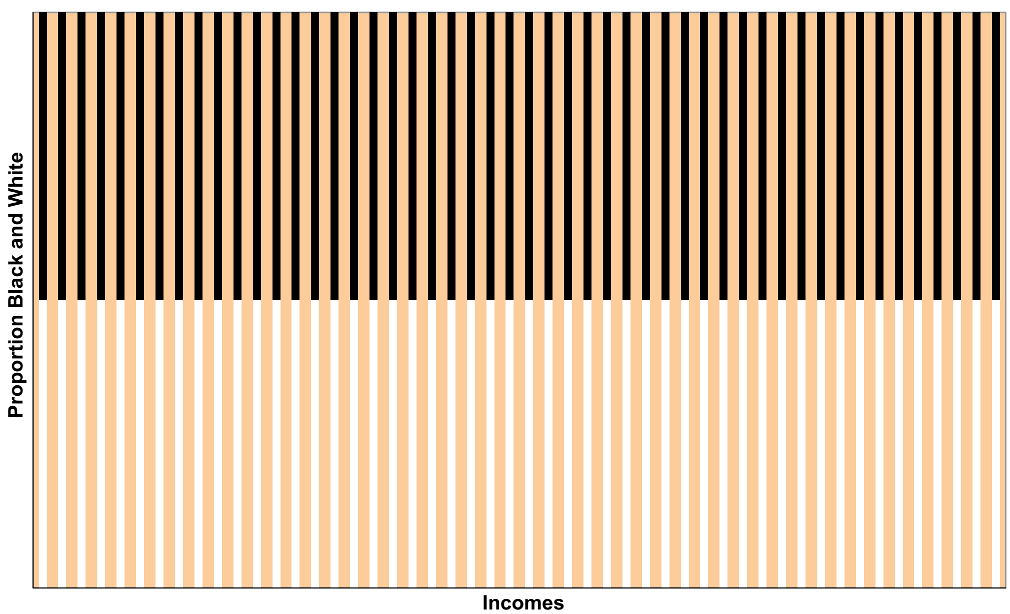

Fig. 1 Zero representational inequality

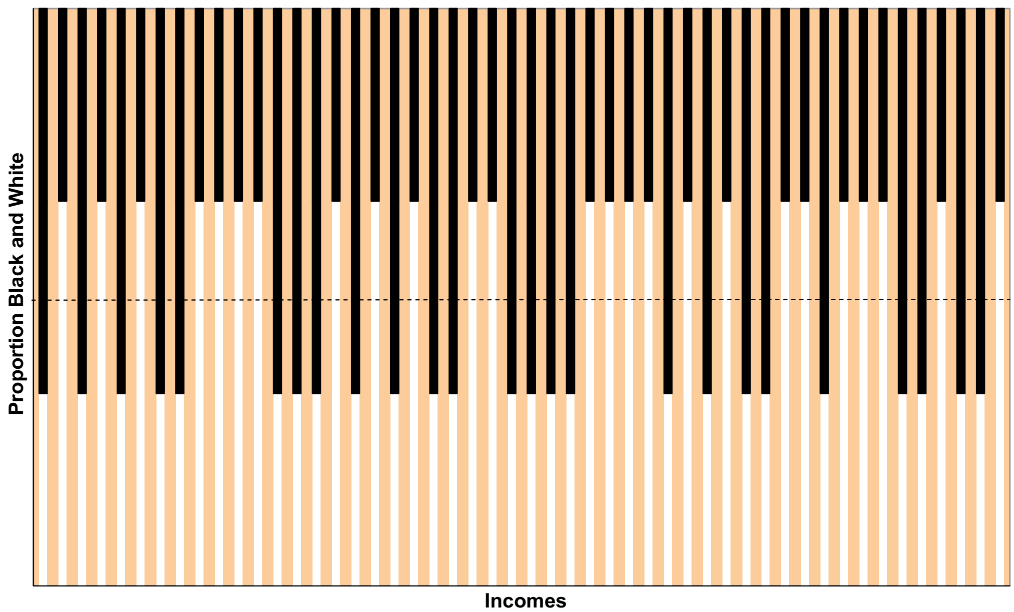

Fig. 2 Nonzero representational inequality

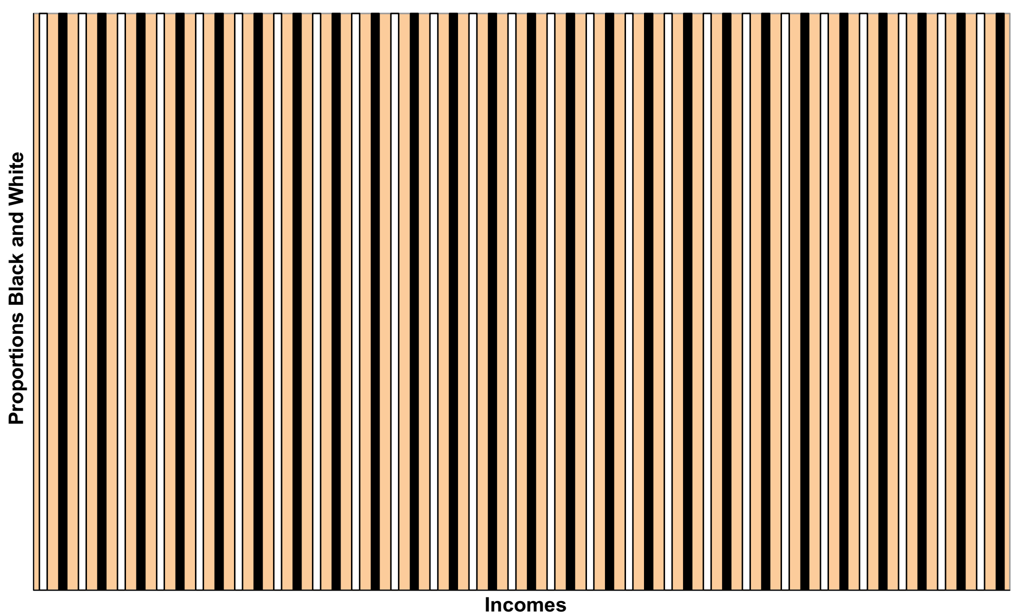

Fig. 3 "Complete" segregation

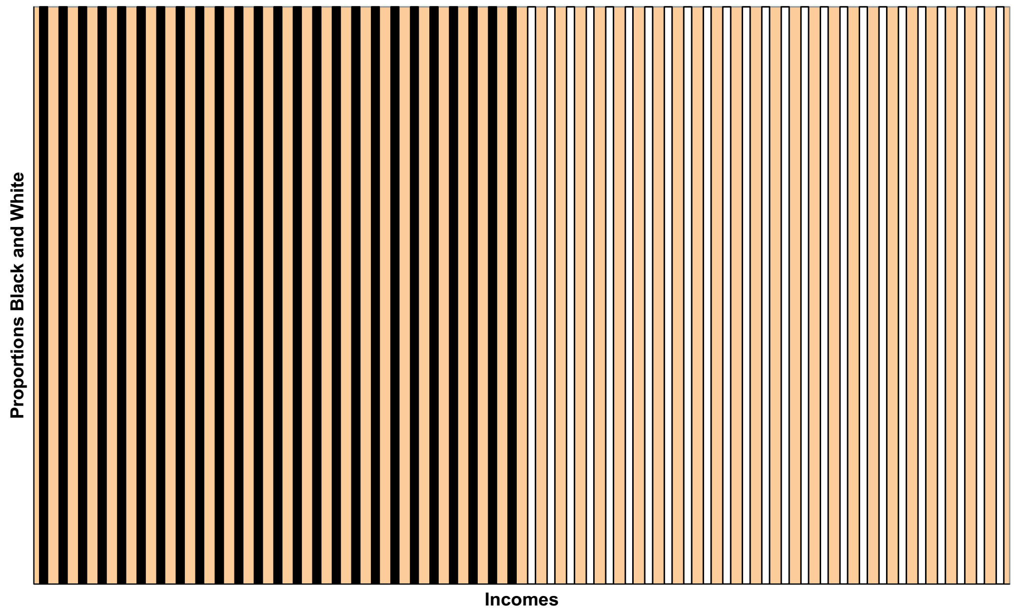

Fig. 4 Polarization

the figures we have discussed here would depict the degree of inequality in political representation.

2.2 Sequence Inequality

The distinction between "complete segregation" and "complete clustering" can be seen by comparing Figs. 3 and 4. Figure 4 depicts the situation that results from a transfer of incomes such that all the whites move to the richer half of society while all the blacks move to the poorer half of society. This situation is one in which each subgroup is concentrated in a different part of the income distribution. Such a situation can plausibly be described as one of "complete clustering" of groups.[7] In both cases, there is complete segregation and thus maximal representational inequality. However, in Fig. 3, whether an individual is black or white provides very little information on his or her rank in society. By contrast, in Fig. 4, whether an individual is black or white provides a great deal of information. One simple way to capture the distinction between Figs. 3and 4 is through the concept of sequence inequality, which together with representational inequality captures the clustering of the income distribution. This concept is linked to the position in the overall societal ranking possessed by individuals belonging to distinct groups in the hierarchy.

An individual (weakly) rank dominates another if that individual is ranked equal to or higher than the other in the possession of the attribute. For any population

[7]Massey and Denton (1988) refer to equivalent concepts.

partitioned into given identity groups, there are a fixed number of between-group pair-wise comparisons between individuals from different identity groups. The share of the total number of such between-group pair-wise comparisons involving a given group in which a member of the group rank dominates a member of some other group is called its level of group rank dominance. Group rank dominance is an indicator of the position the group occupies in the ordinal hierarchy of attribute levels. Another way to understand the difference between Figs. 3 and 4 is simply that the average rank of the whites and the blacks is different. This is clearly a necessary condition for distinct groups to be clustered in different parts of the attribute space. We establish in Appendix One that a monotonic relationship exists between the concepts of group rank dominance and of average rank. Both could be seen to be indicators of the placement of groups in the attribute hierarchy (in the extreme complete clustering of groups) and will thus be referred to as indicators of a group's rank sequence position.

The level of inequality in different groups' rank sequence position (whether as measured by group rank dominance or by average rank) indicates the extent to which a population is clustered. We refer to this concept of inequality as Sequence Inequality (*SI*). Some reflection will suffice to show that this is an unambiguous criterion even when group sizes differ. In any situation sequence, inequality is minimal when the groups are evenly interspersed or symmetrically placed around the median member(s).

It is clear from this discussion that while Figs. 3and 4 depict two groups with equal representational inequality, the two groups possess different levels of group rank dominance and average rank. In Fig. 4, whites have 100% of the available instances of rank domination and higher average rank.

While sequence inequality and representational inequality are related, they are also distinct concepts. A simple example which makes this distinction transparent is provided in Figs. 5 and 6. In Fig. 5, both groups possess the same level of group rank dominance and average rank. The black group has two of the possible four instances of rank domination as does the white group, and their average rank is the same. Thus, there is no sequence inequality between the groups. In the second, both groups again share equally in levels of group rank domination (both have two of the potential four instances once again) and have the same average rank. The situation once again is one in which there is no sequence inequality. However, in the first case, there is complete representational inequality and in the second case there is zero representational inequality. In neither case is group membership always associated with higher rank, yet the cases differ in the degree to which income levels are shared by members of distinct groups.

2.3 Group Inequality Comparison

Figure 4depicts a situation of maximal representational inequality and maximal sequence inequality. It could perhaps be thought of as a situation of polarization in the sense that each group is concentrated at a given pole of the income distribution.

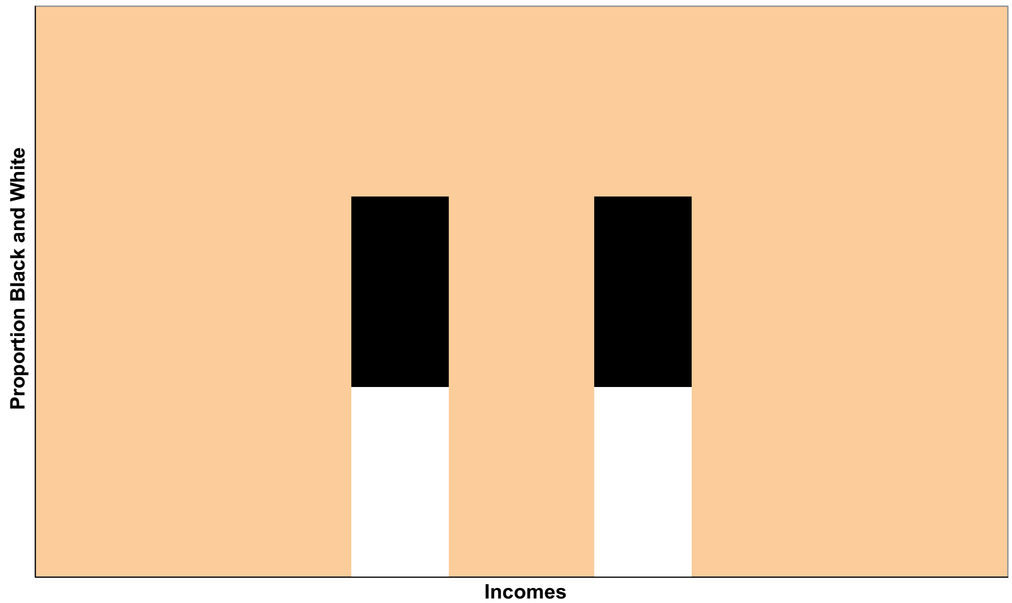

Fig. 5 Perfect sequence equality with perfect representational equality

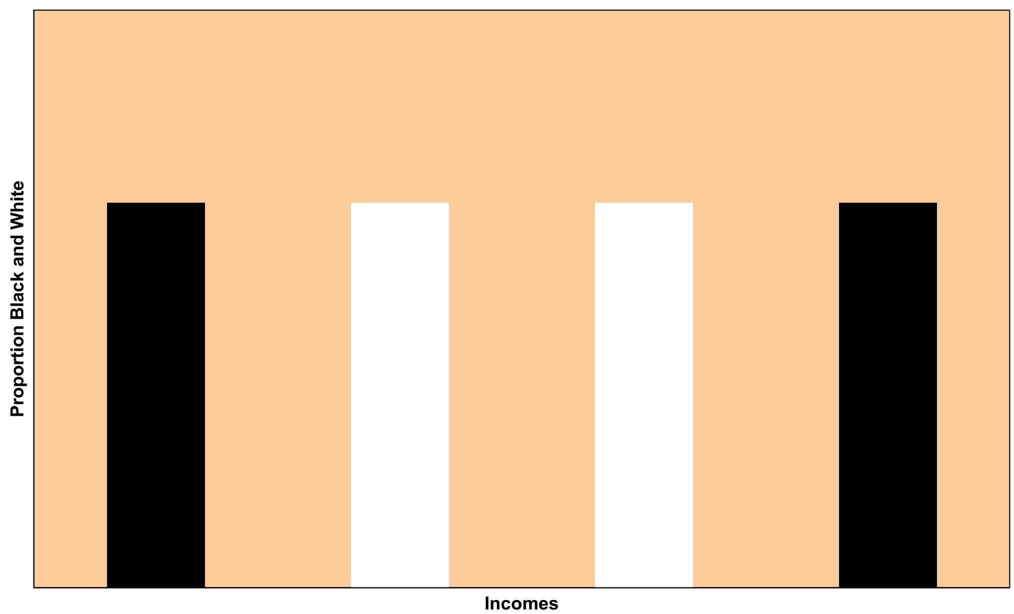

Fig. 6 Perfect sequence equality with complete segregation

However, this is true only in an ordinal sense. Both the situations depicted in Figs. 4 and 7 are *identical* from the standpoints of representational inequality and sequence inequality since neither concept takes note of cardinal information, which alone accounts for the difference between the two situations described. To take account of cardinal information (for instance, concerning the distance between distinct clusters), it is necessary to introduce an additional concept.

A common way to account for such information is to take note of the distance between the means of distinct subpopulations, for example, by using measures of inequality between group means. This indeed is the conception behind *Group Inequality Comparison (II).*However, such an approach ignores relevant information on *within*-group inequality. Consider a two-group society in which all members of each group originally, respectively, possess the mean incomes of their groups. Suppose that both groups experience within-group transfers leading to intragroup inequality. The extent of inequality in the society must be judged to have increased if the measure of inequality employed obeys the Pigou–Dalton Transfer Principle (ensuring that such transfers between persons are deemed to increase overall inequality). However, between-group inequality (understood in terms of inequality between mean incomes of groups) is unchanged. Between-group inequality must be deemed to have become relatively less substantial in comparison with total interpersonal inequality.

An approach to intergroup inequalities which is based on between-group inequalities in isolation rather than on the contribution of between-group inequality to overall interpersonal inequality (i.e., *Group Inequality Comparison (II)*) will fail to contrast situations that might be distinguished. Consider Fig. 8 which depicts a two-group

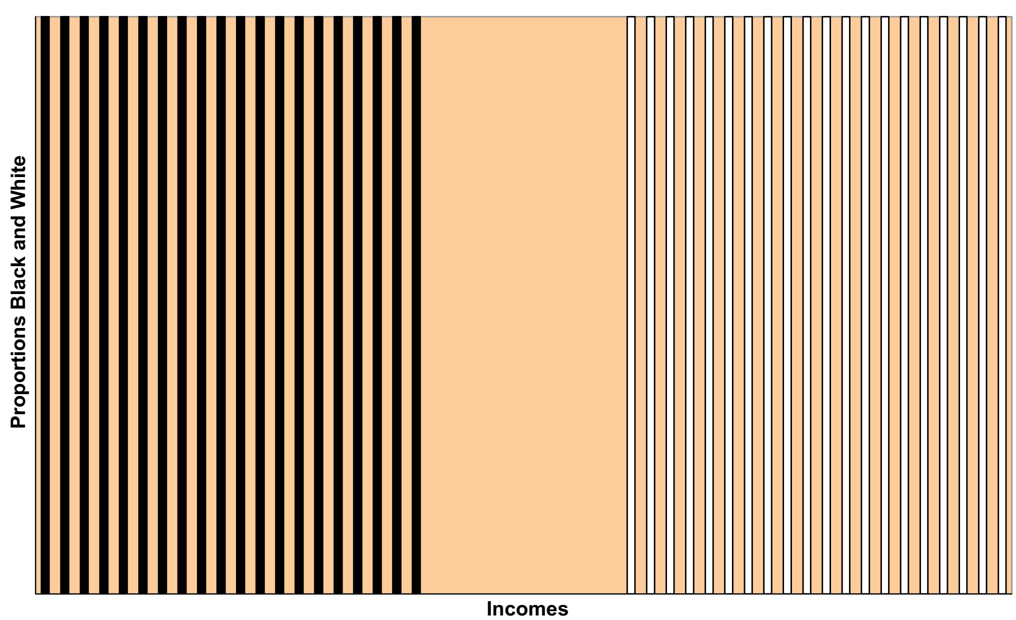

Fig. 7 Polarization

society in which all members of each group originally possess mean income A and B, respectively. Both groups now experience within-group transfers which increase inequality and their distributions are now depicted by densities A' and B', respectively. Assume further that the transfers are such that the span between the means is D and the span between the richest and poorest members of each group is also D. We might plausibly consider intergroup differences to have become less significant after the transfer since no member of the richer group is further away from some member of the poorer group than before the transfer, and all but the very richest member of the richer group is closer to some member of the poorer group.

On the other hand, Group *Inequality Comparison (I)* can have the disadvantage of ignoring information relevant for understanding the extent to which intergroup differences generate overall inequality. To see this, consider what would happen if in Fig. 8, the original populations A and B were made arbitrarily closer to each other while maintaining their separation. According to *Group Inequality Comparison (I)*,there would be no difference between the two situations. If we employed instead the concept of *Group Inequality Comparison (I)* the degree to which between-group differences generate inequality will have fallen. There are potentially good reasons to choose either approach.

Group Inequality Comparison need not be measured, of course, in terms of differences in means and could potentially be understood in other ways—for instance, in terms of differences in medians, generalized means, or other measures of central tendency. Indeed, still other ways of viewing group differences can be envisioned, for example, involving comparison of higher moments of the group-specific distributions of incomes, examination of the extent of "non-overlap" between distributions,

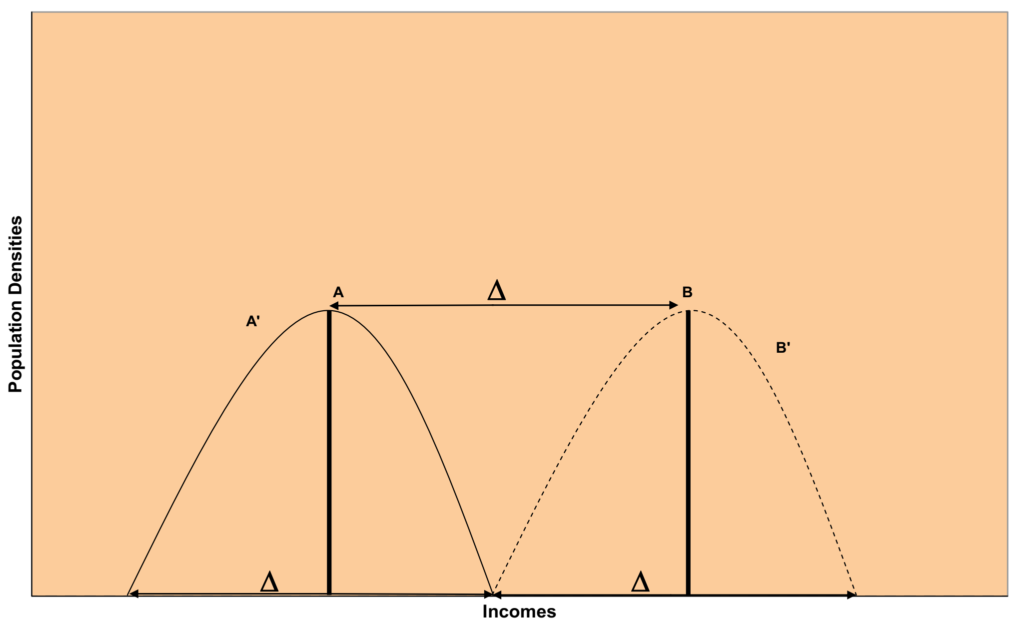

Fig. 8 Group inequality contribution versus inequality between means

etc. For a wide-ranging discussion of methods of defining group separation, see Anderson (2004, 2005). We limit our further discussions of the concept, however, to the case where it is measuring mean differences, for expositional simplicity.

2.4 Combining Concepts: Polarization

We have introduced above three concepts relating to intergroup inequalities: representational inequality, sequence inequality, and group inequality comparison. How are these concepts related to polarization? Polarization is a concept which has been used in many ways in the literature, for example, to mean the absence of "middleness" in a distribution (Wolfson 1994), the distance between the average achievements of groups (Østby, 2008) and the presence of distinct sizable groupings in the income distribution (Esteban and Ray, 1994). Many of these approaches do not explicitly rely on the identification of individuals by identity groups (understood as being distinct from attributes). A contrasting approach understands the level of polarization of a distribution in terms of the extent of intergroup differences in the possession of an attribute. If polarization is defined in this way, it becomes clear that each one of the concepts of intergroup inequality defined above is *itself* a measure of polarization. However, taken individually each may prove to be an unsatisfactory measure of polarization, because of the information to which each is individually indifferent. Thus, the relative ranking of the situations depicted in Figs. 3, 4 and 7 according to the extent of polarization depends on the expansiveness of the approach used. All the figures depict maximal polarization as judged according to *RI*, whereas Figs. 4 and 7 depict maximal polarization according to both *RI* and *SI*, and Fig. 7 depicts more polarization than does Fig. 4 according to *GIC* (taking the figures to possess the same income scale on the horizontal axis).

The fact that our judgments regarding the polarization of society may depend on more than one concept suggests the value of combining measures of intergroup differences to construct orderings of social situations according to the extent of their polarization. Such orderings can be partial and based on dominance of the vectors (two-tuples or three-tuples) defined by the individual measures of intergroup differences, or can be complete if based on some method of aggregation of these measures.

This said, orderings based on combining only a pair of the concepts we have defined (and not all three) will be indifferent to some important considerations that may be deemed relevant in any assessment of polarization. We have already seen that in the two-group case, combining representational inequality and sequence inequality will be sufficient to give us a measure of ordinal polarization. Such a combination however will be indifferent to cardinality and will be unable to distinguish, for example, between the situations depicted in Figs. 4 and 8, respectively.

A measure combining sequence inequality and group inequality comparison is not indifferent to cardinal information on the achievements of individuals but it is indifferent to the degree of clustering of identity groups in any specific income bracket. To see this, consider Figs. 5 and 6 again. Let us assume that, by construction, the mean income of both blacks and whites is the same in both groups in both situations. If this is the case, the index of group inequality comparison is the same in both figures (i.e., zero) and sequence inequality is the same, but representational inequality is different. We may argue that in Fig. 5 there is no clustering of identity groups in distinct parts of the income spectrum, as there is no representational inequality. In Fig. 6, however, blacks are clustered at the top and bottom ends of the income spectrum, and indeed there is complete segregation between the two groups. Note further that we could increase the distance between the blacks at the ends and the whites in the middle, keeping the means of both groups the same (so that the blacks at each end are very distant from the whites at the center) and yet record the same level of polarization defined according to such a measure.

Finally, combining representational inequality and group inequality comparison (I) alone leads to an approach that is indifferent to the sequencing of individuals from distinct identity groups in the income spectrum. Consider the distinction between Fig. 9a and b. Both depict cases of complete segregation. However, in Fig. 9b, some population of blacks has been moved to a higher income than all the whites, thereby increasing within-group inequality for the blacks and total interpersonal inequality. We can further imagine that every white has been given a higher income in such a way that within-group inequality among whites is unchanged and the ratio of between-group inequality to total inequality (which would otherwise have fallen) is restored to its level prior to the initial movement of blacks. In other words, the index of group inequality comparison (I) remains the same by construction, as does representational inequality. However, the sequencing of blacks and whites in the income distribution (and thus sequence inequality) is different. An analogous argument can be made for group inequality comparison (II) by moving the blacks and whites to keep mean incomes of the groups the same.

Any approach to polarization based on a pair of the group inequality concepts we have defined will capture certain judgments about social situations and neglect others. Only by combining all three concepts can an approach to polarization which takes account of the considerations reflected in each of the concepts be constructed.

A variant of group inequality comparison (I) has been proposed as a stand-alone measure of polarization (Zhang and Kanbur, 2001). However, such a measure, while attractive in its simplicity can violate some intuitions. Consider Fig. 10 in which two completely segregated and clustered groups A and B experience within-group progressive transfers which reduce within-group inequality. Further, suppose that they also experience a reduction of between-group inequality through progressive transfers between the members of the two groups in such a way that the ratio of between-group inequality to overall inequality remains unchanged and the groups (whose densities are now depicted by A' and B') overlap. If we utilize group inequality comparison (I) alone as our measure of polarization, a social configuration with A and B is viewed as being exactly as polarized as a situation with A' and B', which seems to conflict with our intuitions. If we, however, combine it with some measure of sequence inequality and/or representational inequality (both of which are lower

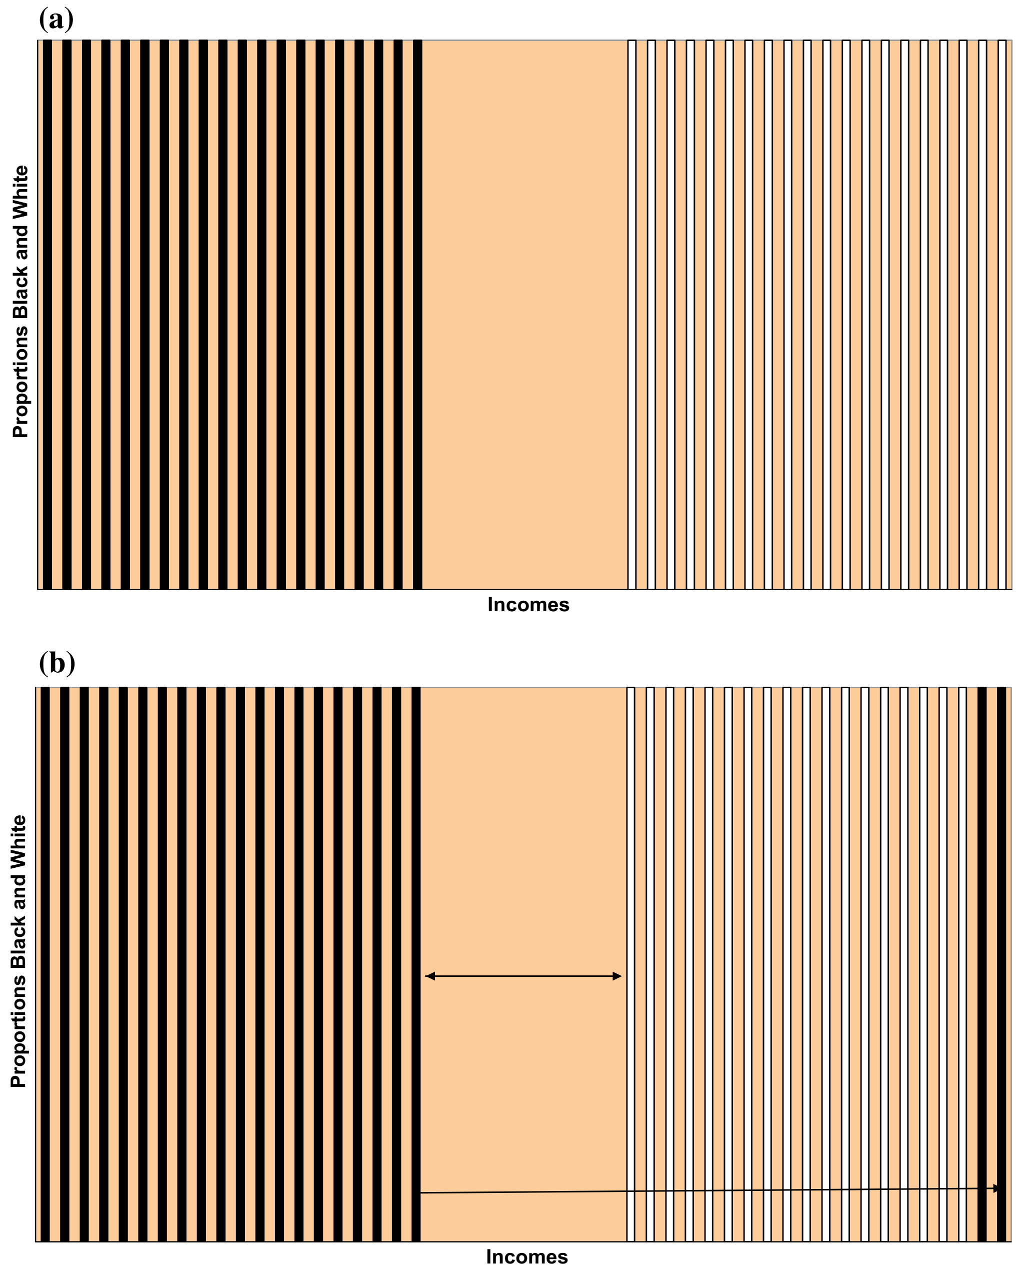

Fig. 9 **a** Maximal representational inequality, maximal sequential inequality with a fixed group inequality contribution, **b** maximal representational inequality, reduced sequential inequality with a fixed group inequality contribution

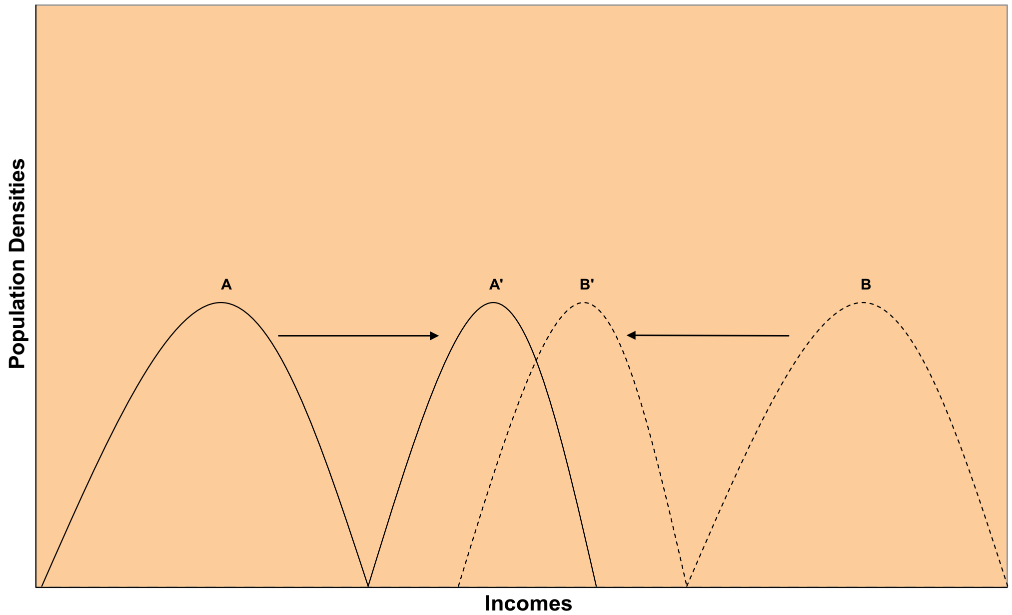

Fig. 10 Group inequality contribution alone is an incomplete measure of polarization

when the groups overlap), the first situation is unambiguously more polarized than the second.

It should be noted that the regressive transfers considered above led to a decrease in the index of group inequality comparison (I), and therefore their impact was in the *opposite* direction from that which would normally be expected of an inequality measure (i.e., to obey the Pigou–Dalton principle of responding to a regressive transfer with an increase in measured inequality). It follows that any measure of polarization which increases when the index of group inequality comparison (I) increases would similarly potentially violate the Pigou–Dalton principle.[8]

3 Part II: From Concepts to Measures

3.1 Formalizing Concepts

Our purpose in this section is to formalize the concepts relating to group differences which we have introduced above and develop measures of them.[9]

[8]This view corresponds to the findings of Esteban and Ray (1994) among others that polarization and inequality are distinct concepts and that measures of polarization need not therefore be expected to obey the Pigou–Dalton principle.

[9]These measures can be readily implemented using a Stata module that we have developed. For an example involving actual data, see Reddy and Jayadev (2011).

We begin by supposing a "social configuration" (ζ) in which there is a population, S_0, of individuals $\{i\}$ of size N partitioned[10] into K distinct identity groups (S_1, $S_2 \ldots S_K$). The individuals possess an attribute (let us say y), drawn from an attribute set, Y. The attributes are not necessarily ordered. For example, the attribute may be a level of income (ordered and cardinally measured), a quality of health (ordered but not cardinally measured), or a club to which a person may belong (distinguished from one another, but not ordered). We employ a superscript to distinguish the information associated with distinct social configurations. For simplicity, we assume (although nothing depends on this other than notation) that the number of elements, l, in the set Y is finite.

More specifically, the individuals $\{i\}$ each belong to a distinct identity group $S_J \subseteq S_0 \quad (J \neq 0)$ so that

$\forall i \in S_0, i \in S_J$ for some $J (J \neq 0)$, with

$$S_J \bigcap S_M = \phi \, \forall J, M \in (1, \ldots K) \text{s.t } J \neq M \text{ and } \bigcup_{J=1}^{K} S_J = S_0.$$

Our assumptions imply there are at least two identity groups which are each smaller than the population as a whole and non-empty. Let the number of persons in group J be denoted by n_J. The proportion of persons of a group J in the society is defined by

$$\theta_J = \frac{n_J}{N} \text{for} \, [J \in (0, \ldots K)].$$

Each individual i has attribute y_i. The same attribute may be shared by more than one individual.

Define the membership function for group $J$by $M_J(y) = \#\{i \in S_J | y_i = y\}$, for $[J \in (0, \ldots K)]$. Moreover, define the complementary membership function for group J by $M_J^-(y) = \#\{i \notin S_J | y_i = y\}$. In other words, the membership function specifies the number of persons in group J who possess attribute y while the complementary membership function specifies the number of persons not in group J who possess attribute y.

[10]We do not consider currently the case of societies in which individuals belong to more than one identity group simultaneously and in which the identity groups do not form a partition of the society into mutually exclusive categories. Generally, a "maximal" partition of a society, generating a mutually exclusive and exhaustive set of groups, can be constructed by generating the Cartesian product of all of the identity groups in the society. This approach may not be deemed appropriate, however, in every situation. For example, a mixed-race group in a society otherwise divided into two races may be deemed to belong to *both* of the races rather than to neither, generating a different characterization of intergroup differences.

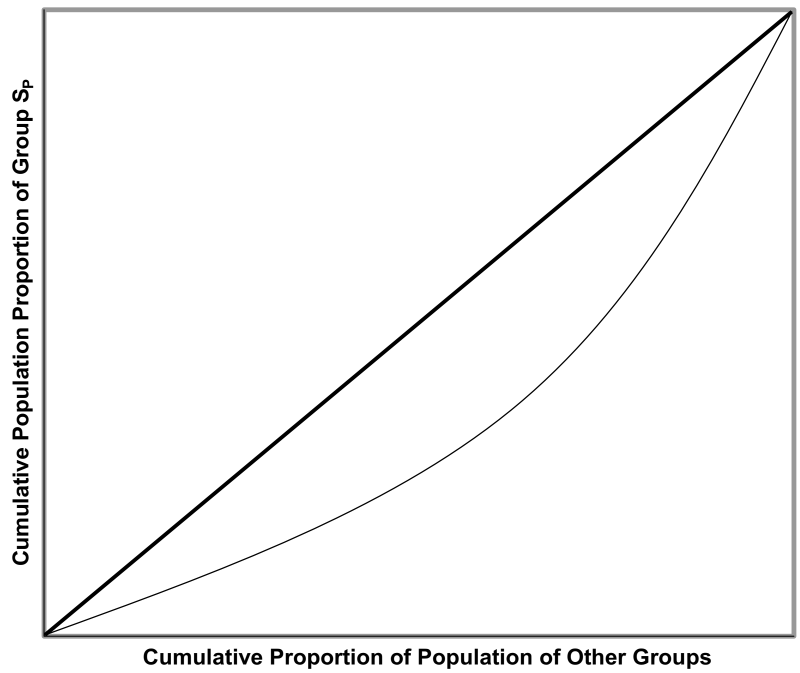

Fig. 11 The representational inequality Lorenz curve

3.2 *Representational Inequality*

A simple way to capture the degree to which each identity group is disproportionately represented among those who share a given attribute would be to describe the ratio of the number of the persons possessing a given attribute who belong to each group, *J*, to their overall number in society for any given attribute (*y*): $\frac{M_J(y)}{M_0(y)} = F_J(y)$. In other words, $F_J(y)$ refers to the proportion of persons who possess a given attribute who belong to group *J*. This information can be captured in what we call the Representational Inequality (RI) Lorenz curve (Fig. 11). As we shall see, this framework allows for a simple way of presenting information concerning these proportions and for analyzing this information using familiar tools.[11]

To construct the *RI* Lorenz Curve for each group, *J*, we first create a rank ordering, R_J, such that

$$F_J(y_{1_J}) \leq F_J(y_{2_J}) \leq \ldots F_J(y_{l_J}),$$

[11] In spirit, this approach is similar to that adopted by Duncan and Duncan (1955) and later, *inter alia*, by Silber (1989, 1991, 1992), and Hutchens (1991, 2004). Other references include Flückiger and Silber (1994). Boisso et al., (1994), and Reardon and Firebaugh, (2002). Silber notes that various information structures (for example, involving the frequencies with which distinct groups possess an attribute such as membership in an occupation) can be analyzed using "measures of dissimilarity" which are analogous to measures of inequality. Our approach builds upon this insight but differs from all of the authors above in explicitly going beyond the two-group case and aggregating information derived from the concentration curves of different groups.

where $F_J(y_{1_J}) \leq F_J(y_{2_J}) \leq \ldots F_J(y_{l_J})$ reflects the ordering of the attributes according to the proportion of the population in the attributes belonging to group *J*.The ordering starts from the attribute for which the proportion of the population consisting of members of group *J* is the lowest and proceeds to the attribute for which the proportion of the population consisting of members of group is the highest.

Clearly, in the case in which the attribute can itself be ordered (e.g., income), the sequence in which the y_{i_J} appear in the ordering R_J will not necessarily be from lowest to highest.

Define:

$\alpha_J(t) = \frac{\sum_{i=1}^{t} M_J^-(y_{i_J})}{N-n_J}$ and $\beta_J(t) = \frac{\sum_{i=1}^{t} M_J(y_{i_J})}{n_J}$, where $[t \in (0, \ldots l)]$ and $\alpha_J(0) \equiv 0$ and $\beta_J(0) \equiv 0$.

The RI Lorenz Curve for group J,$\widehat{L_J}$, can be defined by the following rule, which creates a piecewise linear curve:

When $x = \alpha_J(t)$, for integer values $[t \in (0, \ldots l)]$, then $\hat{L}_J(x) = \beta_J(t)$ and, when x is such that $\alpha_J(t) < x < \alpha_J(t+1)$, $[t \leq (l-1)]$, then

$\widehat{L_J}(x) = \widehat{L_J}(\alpha_J(t))\lambda + \widehat{L_J}(\alpha_J(t+1))(1-\lambda)$, where $\lambda = \frac{x-\alpha_J(t))}{\alpha_J(t+1)-\alpha_J(t)}$.

In using this definition, we follow the procedure described by Shorrocks (1983), p. 5.

This gives rise to a curve as shown in Fig. 11. By construction, the *RI* Lorenz curve must, in the familiar way, begin at (0,0) and end at (1,1), as well as slope upward, with the slope increasing as one moves to the right, since each addition to the total cumulative population of others is associated with an addition of a larger proportion of group *J*. Note that the 45-degree line here has the interpretation of being the line of equiproportionate representation (analogous to the line of perfect equality in the case of an ordinary Lorenz curve). That is, all along this line, the members of identity group *J* are represented at every attribute in the same proportion as they are represented in the population.

Any deviation from the line of equiproportionality represents a situation in which members of the group are disproportionately represented in the possession of certain attributes, leading them to be "over-represented" in the possession of certain attributes and "under-represented" in the possession of others. The *RI* Lorenz curve therefore contains information on the extent of segregation of a population in relation to the attributes possessed. Having defined it, we can draw on the analogy between the *RI* Lorenz curve and the ordinary Lorenz curve to suggest further useful concepts.

Consider for instance what might correspond to the familiar idea of a progressive transfer. Just as a progressive transfer in an income distribution involves a transfer from a person with higher income to a person with lower income, in the context of representational inequality, a progressive transfer could be defined as a transfer of a person from the set of persons who possess an attribute in which his or her identity group is represented more to one in which it is represented less. However, since we are dealing with proportions of identity groups possessing different attributes, a transfer of a single person will change the overall population that possesses each attribute involved, affecting the "denominator" used to assess population proportions for the groups possessing these attributes. We overcome this problem and maintain an unchanged denominator by instead employing the concept of a "balanced bilateral population transfer"[12]:

Definition Balanced Bilateral Population Transfers
Suppose $\exists(y_i, y_j) \in Y$ and S_P, S_Q such that

$$F_P(y_i) > F_P(y_j) \text{ and } F_Q(y_i) < F_Q(y_j)$$

with $P \neq Q$ *and* $i \neq j$.

Then, a progressive (regressive) balanced bilateral population transfer is one in which population mass Δ (i.e., some number of persons; we abstract from integer problems here) of group P is shifted from y_i to y_j and equal population mass Δ of group Q is shifted from y_j to y_i, thereby lowering (raising) $F_P(y_i)$ and $F_Q(y_j)$ while raising (lowering) $F_P(y_j)$ and $F_Q(y_i)$.

A balanced bilateral progressive population transfer results in two upward shifts in the *RI* Lorenz curves for the identity groups (and corresponding downward shifts for regressive transfers). An example of the latter is provided in Fig. 12a and b. The *RI* Lorenz curve that results from a progressive (regressive) balanced population transfer dominates (is dominated by) the *RI* Lorenz curve that preceded the transfer.[13] We note further that:

Lemma 1 *There exists a pair of identity groups and a pair of attributes* (y_i, y_j) *for which a progressive balanced bilateral population transfer can take place if all groups are not equiproportionately represented in the possession of every attribute.*

Proof: See Appendix Three

An *RI* Lorenz curve $\widehat{L}(x)$ weakly dominates an *RI* Lorenz curve $\widehat{L'}(x)$ if and only if $\widehat{L}(x) \geq \widehat{L'}(x)$ for all $x \in [0, 1]$. An implication of this framework is that any Lorenz consistent measure of inequality, for which inequality never decreases when $\widehat{L}(x)$ is replaced by *L(x),* i.e., all income inequality measures used in practice can also be applied to measure representational inequality. It is also well known in the literature on income distribution that it is possible to shift from an income distribution that possesses a Lorenz curve $L(x)$ to another that possesses the Lorenz curve $L'(x)$ where $L(x) \leq L'(x)$ if and only if there exists a corresponding sequence of

[12]This concept of a balanced bilateral population transfer is related to that of a "disequalizing movement" between groups used by Hutchens (2004) in his discussion of a two-group case. However, the latter concept is insufficient in a multigroup case and necessitates the use of the alternative concept which we develop and employ. The concept is also intimately related to the idea of a "marginal preserving swap" which has appeared in the statistical literature (see, for example, Tchen, 1980, Schweizer and Wolff, 1981, and Bartolucci et al., 2001). However, to the best of our knowledge, no one has shown that in the absence of perfect representational equality, there always exists the possibility of achieving a balanced bilateral transfer (as we do in Appendix two) and that this can be given a natural interpretation in terms of Lorenz curves.

[13]For the relevant reasoning, see Shorrocks (1983).

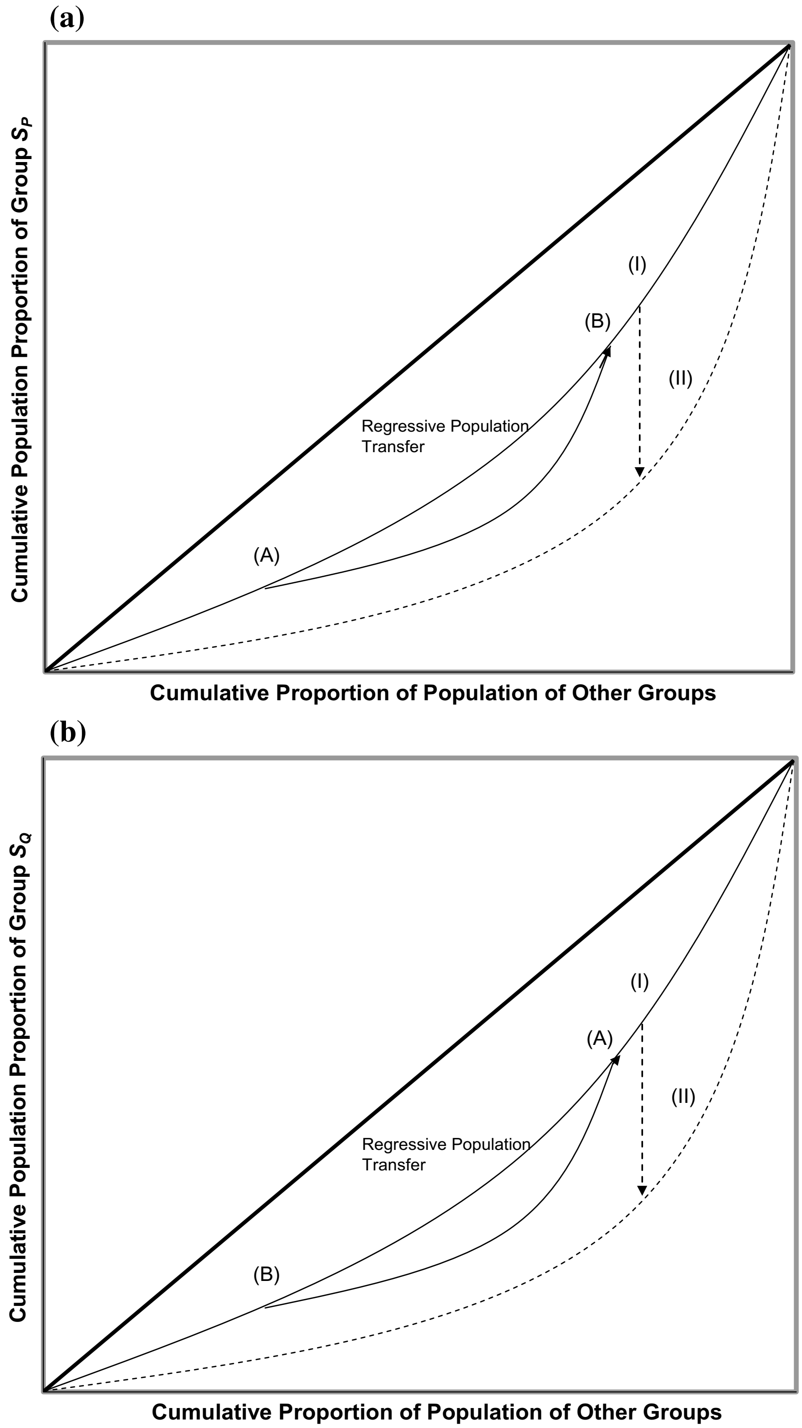

Fig. 12 **a** Balanced bilateral transfer, **b** balanced bilateral transfer

Table 1 Correspondences between conventional inequality and representational inequality concepts

Conventional inequality concept	Representational inequality concept
Inequality	Over- or under-representation
Pigou–Dalton transfers	Balanced bilateral population transfers
(First-order) Lorenz dominance	(First-order) *RI* Lorenz dominance

progressive transfers. Equivalently, in our case, it is possible to shift from a situation for which each group possesses a Lorenz curve $L_J(x)$ to another in which each group possesses a Lorenz curve $L'_J(x)$, where $L_J(x) \leq L'_J(x)$ if and only if there exists a corresponding sequence of balanced bilateral progressive population transfers. For this reason, a balanced bilateral progressive population transfer can be deemed to decrease overall representational inequality.

The consequence is a striking parallel between inequality measures in the income space and inequality measures in the representation space. Table 1 provides a map of the isomorphism between corresponding concepts introduced so far.[14]

Suppose that we apply Lorenz consistent inequality measure $\hat{I}(\hat{L}_J(x)$ to assess representational inequality for group J and denote the resulting vector of measured inequality for all groups in the society by $\bar{\hat{I}}$ and its individual components by $\bar{\hat{I}}_J = \hat{I}(\hat{L}_J(x))$, $J \in (1, \ldots K)$. Then, an overall measure of representational inequality in the society is given by $RI = f(\bar{\hat{I}}, \ldots)$, where $f(\bar{0}) = 0$, $f(\bar{1}) = 1$, and $\frac{\partial f}{\partial \hat{I}_J} \geq 0$ for all $J \in (1, \ldots K)$. One simple version of such an aggregation function, *f*, is the mean of the group-specific representational inequality measures. It may seem attractive for a measure of overall representational inequality to take into account subgroup sizes and respond to unequally sized groups differently. Indeed, it will be argued below that there can be sound reason for such weighting. We may define a population-weighted overall representational inequality measure of the following form:

$$RI = \frac{1}{K}\sum_{J=1}^{K}(\theta_J)\Big(\hat{I}(\hat{L}_J(x))\Big),$$

where θ_J refers to the population weight of subgroup *J*.

Such a measure can be offered some justification through axiomatic underpinnings which we consider in the next section.

[14]The concepts of the generalized Lorenz curve and dominance of generalized Lorenz curves do not possess straightforward and interpretatively useful analogs in the area of representational inequality since the concept of an income mean does not possess a straightforward analog in this realm.

3.3 Sequence Inequality

As noted in the discussion of the previous section, representational inequality is a measure of group differences which is indifferent to the ordering of attributes as well as to their cardinal properties. To operationalize our concept of sequence inequality, therefore, we now assume that the attributes can be ordered.[15]

Considering first the concept of group rank dominance, we define a pair-wise individual rank domination function, δ_{ij}, for a given pair of individuals i and j as follows: $\delta_{ij} = 0$ if $y_i < y_j$ and $\delta_{ij} = 1$ if $y_i \geq y_j$. We can now define the group rank domination quotient for group J as follows: $\tau_J = \frac{\sum_{i \in K} \sum_{j \in (J \neq K)} \delta_{ij}}{n_J(N-n_J)}$. It can be seen that τ_J possesses the interpretation of the proportion of possible instances of pair-wise domination involving members of group J and members of other groups in which such domination actually occurs. It is evident that this quotient varies between a minimum of 0 and a maximum of 1 for any group. The size of the group plays no direct role in determining the value of the group rank domination quotient. Rather, it is the placement of members of the group relative to members of other groups that determine the quotient. Sequence inequality could be treated simply as the measured inequality in τ_J across groups. It is common for individuals from a given group to express pride or shame at the achievements or failures of other members from that group. Such a psychological interpretation can provide justification for treating the group rank domination quotient as defining the experience of everyone in that group and measure inequality across all individuals in possession of that experience.[16]

A seemingly puzzling asymmetry is implied by our approach to sequence inequality. Consider two populations, consisting of one white individual and ten black individuals each. In the first population, the individuals are ordered in the income space from lowest to highest as (w, b, b, b,….b), and in the second, the individuals are ordered from lowest to highest as (b, b, b, b…..w). In the first instance, all 10 black members possess a domination quotient of 1, while the white individual possesses a domination quotient of zero. The inequality in domination quotient is therefore inequality in a population having scores (0, 1,1,1,1….1). In the second case, all 10 black members possess a domination quotient of 0, while the white individual possesses a domination quotient of 1. The inequality in domination quotients is therefore the inequality in a population having scores (0,0,0,0,0….1). More sequence inequality will be recorded in the first case than in the second, even though all that has been done is to change the placement of the white from being at the bottom to being at the top of the income spectrum. While this may initially appear puzzling, it is perhaps appropriate to treat these cases asymmetrically. By the psychological interpretation,

[15] There is a small nascent literature on the measurement of ordinal inequality. Some key references include Allison and Foster (2004), Reardon (2008) and Abul Naga and Yalcin (2009).

[16] One way to interpret sequence inequalities is in terms of an analogy to a society wherein each group practices radical egalitarianism. In such a society, an even distribution of each group's share of the social assets, in this case instances of rank domination, results among the individuals belonging to the group.

in the first instance, most people in society do not experience a relative deprivation. By contrast, in the second, most do.

As we noted above, the average rank of a group (call it $\omega_J,\ J \in (1, \ldots K)$) is also an indicator of group rank sequence position. In fact, it is linked in a direct and monotonic fashion to group rank dominance. It is easily shown that the relation between them, for a perfectly segregated population is

$$\tau_J = \frac{\omega_J - (n_J + 1)/2}{N - n_J}.$$

In the case of populations which are not perfectly segregated, appropriate changes to the definition of a rank maintain this relationship (see Appendix Two):

The inequality in group rank sequence position across groups can be assessed either in terms of the inequality of group rank dominance quotients or that of average group ranks.[17] In either case, if a member of a group (the "beneficiary") exchanges his or her attribute with another person in a different group who has a higher level of the attribute, then the indicator of group rank sequence position is increased for the group to which the beneficiary belongs and is decreased for the other group. We assume henceforth in this section that we are specializing to the case of group rank dominance quotients, although the concepts we present can equally be applied to average ranks.

The group rank dominance quotients achieved by members of distinct groups can be captured by what we call the Group Rank Dominance (GRD) Lorenz curve. The GRD Lorenz curve relates the cumulative proportion of the total of the group rank domination quotients to the cumulative population of groups, when the identity groups are ordered from lowest group rank domination quotient to highest. It captures the degree of inequality in group rank domination quotients. Any symmetric arrangement of identity groups in the attribute space (i.e., one in which for any instance in which a member of a given group rank dominates a member of another group, a distinct pair can be found in which the opposite is true) is one of the perfect equalities in group rank domination quotients, and will give rise to a GRD Lorenz curve which is on the 45-degree line.

We can now define coordinates of the GRD Lorenz curve associated with each group added as follows:

$\alpha(t) = \frac{\sum_{J=1}^{t} n_J}{N}$ and $\beta(t) = \frac{\sum_{J=1}^{t} \tau_J}{\sum_{J=1}^{K} \tau_J}$, where $[t \in (0, \ldots K)]$ and $\alpha(0) \equiv 0$ and $\beta(0) \equiv 0$.

We can now define the GRD Lorenz curve as a whole, $\widehat{LGRD}$, as follows.

[17]For a given group, although the ordinal ranking of social configurations according to the group's rank sequence position does not depend on the choice between these indicators (or indeed any other monotonic transformation thereof) the cardinal level of the indicator does depend on it. As a result, the choice of indicator can be consequential for determining the measured sequence inequality.

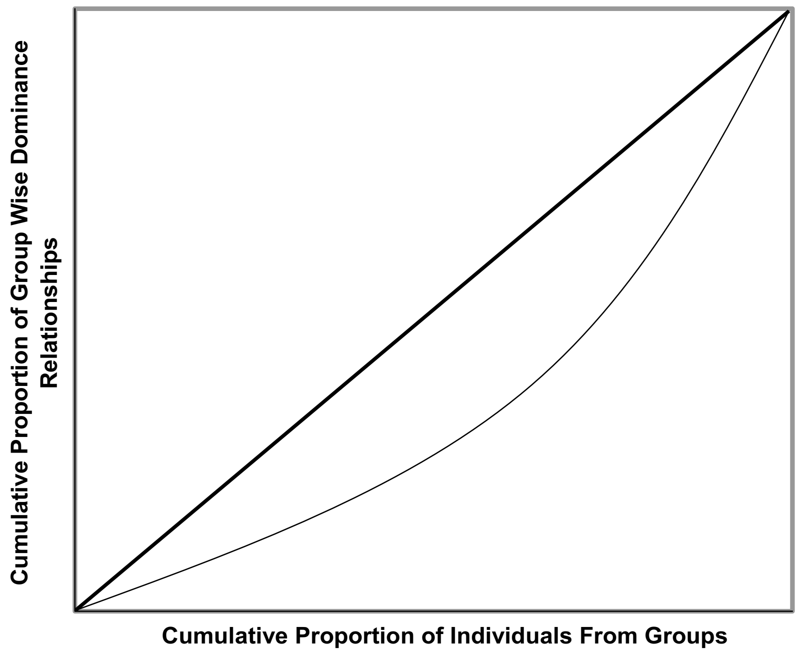

Fig. 13 The rank dominance Lorenz curve

When $x = \alpha(t)$, for integer values $[t \in (0, \ldots K)]$, then $\widehat{LGRD}(x) = \beta(t)$and, when x is such that $\alpha(t) < x < \alpha(t+1)$, $t \leq (K-1)$, then $\widehat{LGRD}(x) = \widehat{LGRD}(\alpha(t))\lambda + \widehat{LGRD}(\alpha(t+1))(1-\lambda)$, where $\lambda = \frac{x-\alpha(t))}{\alpha(t+1)-\alpha(t)}$.

An example of such a curve is shown in Fig. 13. Since the Lorenz curve is defined for sequence inequality analogously to income inequality, with income corresponding to the group rank domination quotient of the groups to which individuals belong, the properties of the GRD Lorenz curve are analogous to those of the ordinary Lorenz curves. Once again, therefore, any Lorenz consistent measure of inequality will suffice to capture the level of sequence inequality.

3.4 Group Inequality Comparison

Group Inequality Comparison (I) refers to the degree to which between-group inequalities contribute to overall inequality. Typically, measures which are "additively separable" (such as members of the generalized entropy class) have been utilized for this purpose (see, for example, Shorrocks 1980, Foster and Shneyerov 1999 and Zhang and Kanbur 2001), although such a restriction is not required. In particular, if the between-group inequality is defined as the inequality that arises when every member of the population is assigned a representative level of an attribute (mean, generalized mean, median, or other measures of central tendency) of the group to which they belong, then the ratio of between-group inequality to total interpersonal inequality can serve as an index of Group Inequality Comparison (I). This measure has the advantage of always lying between zero and one and responding in an appropriate way to intragroup transfers. More generally, any indicator that the distributions associated with different groups are different can potentially serve as a measure of Group Inequality Comparison.

3.5 Polarization

Polarization as we have defined it above aggregates the three concepts concerning group differences which we have defined. The range of polarization measures which could be used is very wide indeed since any such measure could involve any form of aggregation of a three-tuple (*RI*,*SI*, and *GIC*), and in turn each element of this three-tuple could be defined in various ways. Further, any measure of polarization which is positively responsive to all three will only be maximized in a situation where all three are maximized.

An empirical examination which involves these four concepts can, as we have noted, be achieved using almost any common measure of inequality. The choice will naturally bring in additional implications and properties. Given this flexibility, an analyst can choose which measure to utilize to satisfy the additional properties thought important. Thus, for example, a researcher who wishes to treat sequence inequality as being decreased more in a situation where an exchange of ranks happens between members of different groups, each of whom has lower ranks to begin with, can choose an inequality measure which shows the required form of transfer sensitivity (e.g., a generalized entropy index with appropriately chosen parameters). Whether the measure of polarization can be normalized in a specific way will also depend on the choice of the underlying measures of inequality.

4 Part III: Axiomatic Framework

We define below some requirements that may reasonably be imposed on measures of each of the concepts defined above, considering each of them in turn. We also identify some classes of measures which satisfy these requirements.

4.1 Axioms (Representational Inequality)

We begin by suggesting some requirements which may be imposed on an overall representational inequality measure *RI* when it is viewed as a function of the information in a social configuration ζz. We write $RI = RI(\zeta)$ to reflect this dependence.

Axiom (RI1): *Lorenz Consistency*

Let *(ζ^1, ζ^2)* refer to two different social configurations and *(I,J)* refer to two different identity groups. If *(ζ^1, ζ^2)* are such that $\hat{L}_I^1 \geq \hat{L}_I^2$, $\hat{L}_J^1 \geq \hat{L}_J^2$, and $\hat{L}_H^1 = \hat{L}_H^2$, $\forall H \neq I, J, H \in (1, \ldots K)$, then *RI($\zeta^1$) ≤ RI(ζ^2).*

In other words, all else remaining equal, a social configuration which is at least as segregated according to the criterion of Lorenz dominance of representational inequality Lorenz curves is one which is at least as representational unequal. It may be noted that just as there is an equivalence between Lorenz consistency of an inequality measure and that measure's respect for the Pigou–Dalton Transfer Principle, there is an equivalence between Lorenz consistency of a representational inequality measure as defined here and the requirement that the representational inequality measure respond to a progressive balanced bilateral transfer by registering a decrease.

Axiom (RI2): Within-Group Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))_J$ and
if *(ζ^1, ζ^2)* are such that $y_{iJ}^1 = y_{\pi_J(i)J}^2 \forall(i)$, where π_J is a permutation operator applied to $(1, \ldots n_J))$, then *RI(ζ^1) = RI(ζ^2).*

In other words, a measure of overall representational inequality is invariant to permutations of the attributes assigned to individuals within an identity group.

Axiom (RI3): Group Identity Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))$ and if *(ζ^1, ζ^2)* are such that $y_{iJ}^1 = y_{i\pi(J)}^2$ and $n_J^1 = n_{i\pi(J)}^2$, $\forall(i)$,where π is a permutation operator applied to $(1, \ldots K)$, then *RI(ζ^1) = RI(ζ^2).*

In other words, a measure of overall representational inequality is invariant to permutations of the group identities with which distinct sets of individual attributes are associated. This axiom incorporates the idea that all the information relevant to assessing representational inequality is taken into account by noting the partition of the society into groups and the attributes of the members of these groups. The axiom embodies the idea that there is no need to take independent account of any other features of groups. This approach disallows the incorporation of judgments that group identities are *additionally* relevant (e.g., because of past histories or present injustices not already reflected in the information described by the social configuration).[18]

Axiom (RI4): Minimal Representational Inequality

Let $\hat{L}_E$ be the RI Lorenz curve corresponding to even representation (i.e., the line of equiproportionate representation). If $\hat{L}_J = \hat{L}_E \; \forall J \in (1, \ldots K)$, then *RI* =0.

In other words, minimal overall representational inequality is achieved when all identity groups are represented in the same proportion as their share of the population for all attributes, and has measure zero.

---

[18]See Loury, (2004) for an extensive discussion on the merits of the anonymity axiom as applied to groups.

Axiom (RI5): Maximal Representational Inequality

The maximum level of Representational Inequality is 1.

This is a normalization axiom which may be imposed for interpretative convenience. It may be dispensed with if it is desired to employ an unbounded inequality measure (such as a measure of the additively decomposable generalized entropy class).

Axiom (RI6): Positive Population Share Responsiveness of Overall Representational Inequality

Suppose that a measure of overall representational inequality is a function of the vector of measures of representational inequality of groups, $\bar{\hat{I}}$. Suppose further that the population share for group J is increased and that for group H is decreased, and the set of measures of representational inequality of groups remains unchanged as do the population shares for any remaining groups. Suppose further that $\bar{\hat{I}}_J > \bar{\hat{I}}_H$, i.e., that the group-specific representational inequality of group J is greater than that of group H. Then, the measure of overall representational inequality must increase.

This axiom can be motivated in different ways. We might, for example, believe that a group which is very small in the population but which is highly unequally represented simply because it is a small group in a society where there is unequal representation should not affect overall representational inequality in the same manner as a group which is much larger.

We may note that the measure of overall representational inequality defined above, $RI = \frac{1}{K}\sum_{J=1}^{K}(\theta_J)\left(\hat{I}(\hat{L}_J(x))\right)$, satisfies these axioms if the measure used to assess representational inequality for each group, $\hat{I}$, is Lorenz consistent, which will be the case if it has the form of any standard inequality measure, for example, the Gini coefficient.

From another perspective, it may not be appropriate disproportionately to disvalue the unequal representation of smaller groups. If one is interested in the experience of groups as opposed to the experience of individuals within groups, it should make no difference whether the group is small or large. Following this intuition, there is no reason to promote a population-weighted overall measure and one should instead adopt a measure which weights every group equally. This alternative may seem especially compelling if one views polarization as an attribute of the society as opposed to the individuals who belong to it. Such a measure can satisfy all the other axioms.

4.2 Axioms (Sequence Inequality)

In what follows, we shall use γ_J to refer to the indicator of group rank sequence position (which may be either the group rank domination quotient or the average rank) of group J. Let *SI* refer to the measure of overall sequence inequality. Some reasonable axioms are as follows:

Axiom (SI1): *Lorenz Consistency*

Let (ζ^1, ζ^2) refer to two different social configurations and *(I,J)* refer to two different identity groups. Further, let $\hat{L}$ refers to the Lorenz curve describing inequality across groups in the indicator of group rank sequence position, γJ. If (ζ^1, ζ^2) are such that $\hat{L}^1 \geq \hat{L}^2$, then $SI(\zeta^1) \leq SI(\zeta^2)$.

Axiom (SI2): Within-Group Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))$ and
If (ζ^1, ζ^2) are such that $y^1_{iJ} = y^2_{\pi_J(i)J} \forall(i)$, where π_J is a permutation operator applied to $(1, \ldots n_J))$, then $SI(\zeta^1) = SI(\zeta^2)$.

Axiom (SI3): Group Identity Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))$ and if (ζ^1, ζ^2) are such that $y^1_{iJ} = y^2_{i\pi(J)}$ and $n^1_J = n^2_{i\pi(J)}$, $\forall(i)$,where π is a permutation operator applied to $(1, \ldots K)$, then $SI(\zeta^1) = SI(\zeta^2)$.

Axiom (SI4): Sequence Inequality Limits

Let $\hat{L}_E$ be the Lorenz curve (describing inequality in the indicator of group rank sequence position, γ_J) that corresponds to even group rank sequence position (i.e., the case in which γ_J is the same for all groups). If $\hat{L} = \hat{L}_E$, then $SI = 0$.

Axiom (SI5): Maximal Sequence Inequality
The maximum level of sequence inequality is 1. As with Axiom *RI5* above, this is a normalization axiom which may be imposed for interpretative convenience. It may be dispensed with if it is desired to employ an unbounded inequality measure (such as a measure of the additively decomposable generalized entropy class).

4.3 Axioms (Group Inequality Comparison)

Some reasonable axioms may be as follows, if members of each group, *j*, are assigned a representative income, μ_j, and possesses an individual income, y_{ij}.

Axiom (GIC1): *Between-Group Synthetic Population Lorenz Consistency*

Let (ζ^1, ζ^2) refer to two different social configurations. Assume that a synthetic population is constituted in which every member of a group, *j*,is assigned the same representative income for its group, μ_j. Consider the Lorenz curve, $\tilde{L}^1$, $\tilde{L}^2$, for the resulting synthetic population in each social configuration. If (ζ^1, ζ^2) are such that $\tilde{L}^1 \geq \tilde{L}^2$ and $L^1 = L^2$(i.e., the overall Lorenz curves for the actual population remain unchanged), then GIC$(\zeta^1) \leq$ GIC(ζ^2).

This axiom states that between-group regressive transfers which do not change the overall interpersonal distribution must have an appropriate directional effect (nondecreasing) on the measure of GIC. Thus, for example, an exchange of incomes between individuals of different incomes belonging to two different groups that results in an increase in inequality in the synthetic population must increase the measure of GIC.

Axiom (GIC2): Within-Group Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))$ and
if (ζ^1, ζ^2) are such that $y^1_{iJ} = y^2_{\pi_J(i)J} \forall(i)$, where π_J is a permutation operator applied to $(1, \ldots n_J))$, then $GIC(\zeta^1) = GIC(\zeta^2)$.

Axiom (GIC3): Group Identity Anonymity

If y_{ij} represents the attribute of person i ($i \in (1, \ldots n_J)$) belonging to group J $((J \in (1, \ldots K))$ and if (ζ^1, ζ^2) are such that $y^1_{iJ} = y^2_{i\pi(J)}$ and $n^1_J = n^2_{\pi(J)}$, $\forall(i, J)$, where π is a permutation operator applied to $(1, \ldots K)$, then $GIC(\zeta^1) = GIC(\zeta^2)$.

Axiom (GIC4): Within-Group Lorenz Consistency

Let (ζ^1, ζ^2) refer to two different social configurations. Further, let $\hat{L}^1_i$ and $\hat{L}^2_i$ refer to the Lorenz curves describing inequality within each group, i, in the respective social configurations. If (ζ^1, ζ^2) are such that $\hat{L}^1_i \geq \hat{L}^2_i$, but $\mu^1_i = \mu^2_i$, then GIC((ζ^1) $\geq$ GIC((ζ^2).

This axiom states that within-group (weakly) regressive transfers of income must have an appropriate directional effect (nonincreasing) on the measure of GIC, holding the representative incomes of groups constant. Clearly, since *Group Inequality Comparison (II)*does not rely on any information about within-group inequality, imposing this axiom will exclude its use.

It may be readily checked that a measure of *GIC* of the form *B/T*, where *B* represents the inequality measure for the synthetic population in which each member of the society is assigned the representative income of its group and *T* represents the total interpersonal inequality of the society, which satisfies all the axioms above. Such a measure would capture the concept of Group Inequality Comparison (I). In contrast, employing *B* alone as the measure of *GIC* would capture the concept of Group Inequality Comparison (II). Such a measure would satisfy Axioms (*GIC1*) $-$ (*GIC3*) alone.

4.4 Axioms (Polarization)

We have proposed above to define polarization as a function of the other concepts of group difference we have defined. In a working version of this paper (Reddy and Jayadev, 2011), we provide conditions that yield a simple example of a polarization measure that permits the underlying inequality measure used to calculate *RI* and *SI* to be chosen flexibly as long as it is bounded and normalized to vary between zero and one[19]

$$P = (RI)(SI)(GIC).$$

It can be shown that it is the unique measure which satisfies the required conditions and is of the CES functional form. The measure is used to characterize group-based differences in various societies in Jayadev and Reddy (2011).

It is interesting to note that the circumstances in which this measure of polarization is maximized are different from those identified in Duclos, Esteban, and Ray (2004), in which this happens when there are two equal sized groups. The measure of polarization identified here can be maximized regardless of the number of groups, and to approach its maximum it is required that the poorer group be as large as possible relative to the richer groups, that there is complete segregation and that there is no within-group inequality (Fig. 14).

5 Part IV: Conclusion

This paper has sought to clarify how one may assess social situations according to the extent to which attributes are disproportionately possessed by different social groups. The measures we have developed capture the various ways in which experiences of members of distinct groups may differ. Thus, social situations can differ in the extent to which members of a group share experiences with members of other groups (representational inequality), experience the same or different relative positions (sequence inequality) and experience differences in the extent to which interpersonal inequalities are accounted for by intergroup differences (group inequality comparison). These concepts are distinct but complexly interrelated. They each integrate empirical observations and evaluative judgments. Judgments concerning the relative importance to be attached to different aspects of intergroup difference are also involved when they are combined (for example, to form a measure of polarization).[20] These measures have an intuitive appeal and can have widespread application in social science.

There appear to deep-seated tendencies for societies to exhibit segregation, clustering, and polarization of identity groups. This observation has important implications for both empirical investigations of societies and for social evaluation. We hope that, given Satya Chakravarty's lifelong concern with the assessment of social inequalities, he and others may find the concepts and measures that we have discussed to be useful.

[19]The requirement that *RI* and *SI* are bounded and normalized to vary between zero and one excludes certain inequality measures, such as the additively decomposable members of the generalized entropy class.

[20]The concepts we have discussed can be understood as "thick ethical concepts," on which see, e.g., Putnam (2004).

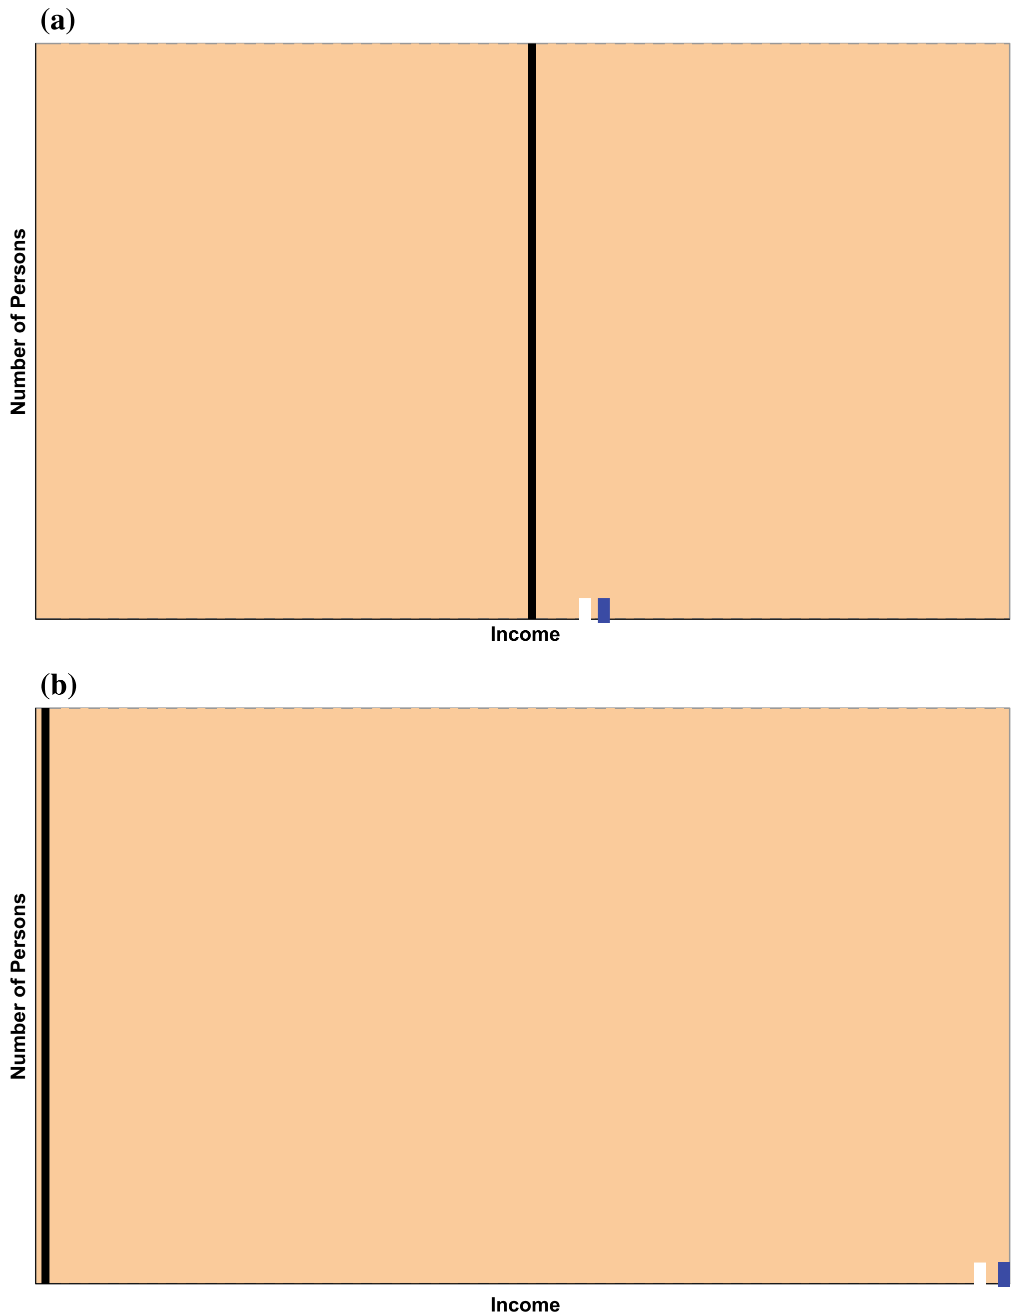

Fig. 14 **a** Very high polarization with GIC (l), **b** Very high polarization with GIC (lI)

Acknowledgements We would like to thank for their useful comments or suggestions (without implicating them in errors and imperfections) James Boyce, Indraneel Dasgupta, Rahul Lahoti, Hwok-Aun Lee, Glenn Loury, Yona Rubinstein, Peter Skott, Rajiv Sethi, Joseph Stiglitz, S. Subramanian, Roberto Veneziani, and other participants at seminars in the Dept. of Economics at Brown University, the University of Massachusetts at Amherst, the Jerome Levy Institute at Bard College, Queen Mary, University of London and the Brooks World Poverty Institute at the University of Manchester, and an anonymous referee for the current volume.

Appendix 1 (Rank Domination Quotient and Average Rank)

As we noted above, the average rank of a group (call it ω_J, $J \in (1, \ldots K)$) is also an indicator of group rank sequence position. In fact, it is linked in a direct and monotonic fashion to group rank dominance. As before, we understand rank as referring to the position in which an individual appears when incomes are sequenced from lowest to highest (the ascending order of values). When individuals from the same group share an income, we shall assign them a rank equal to the average position in which an individual appears when incomes are sequenced from lowest to highest. We shall consider subsequently the rule to be applied in assigning ranks when individuals from different groups share an income.

Consider at the outset, for simplicity, a perfectly segregated population in which there is no more than one individual in each income bracket. In such a population, the total number of instances of pair-wise rank domination that members of group J enjoy vis-à-vis others can be understood as a function of the ranks of members of group J in the population. The lowest ranked member of group J, having rank r_1 dominates $(r_1 - 1)$ persons belonging to other groups. The second lowest ranked member of group J, having rank r_2 dominates $(r_2 - 2)$ persons belonging to other groups (i.e., $(r_2 - 1)$ persons belonging to all groups -1 person belonging to the same group)). Extending this logic, the total number of instances of pair-wise domination by members of group J is

$$\sum_{i=1}^{n_J} (ri - i).$$

The rank domination quotient correspondingly is

$$\tau_J = \frac{\sum_{i=1}^{n_J} (r_i - i)}{N - n_J}$$

from which it follows that:

$$\tau_J = \frac{\omega_J - (n_J + 1)/2}{N - n_J}.$$

It is easy to see that this formula also applies in the case in which there may be more than one person in an income bracket but all persons who share an income bracket are always from the same group. In contrast, in the most general case of populations in which there may exist some income brackets which contain members of distinct groups there can be ties in the income ranks assigned to members of different groups, which will imply that this formula will no longer hold exactly unless the ranks are assigned appropriately to individuals in the same income bracket. Specifically, if strict domination is the concept that is employed then this relationship will hold exactly if individuals in the same income bracket are assigned a rank equal to the lowest of their positions in the ascending order of values. Correspondingly, if weak domination is the concept that is employed, then this relationship will hold exactly if each individual in the income bracket is assigned a value equal to the sum of the lowest of the positions of the individuals sharing the income bracket (in the ascending order of values) and the number of individuals from other groups with whom they share the income bracket.

The correspondence we have derived between τ_J and ω_J holds also in the case of continuous distributions, as can be shown through limit properties. In this case, the average rank of members of a group, *J*, is defined by

$$\omega_J = N \int F(x) g_J(x) dx$$

and the rank domination quotient for the group is defined by

$$\tau_J = \int (F(x) - \theta_J g_J(x)) dx,$$

where $F(x)$ is the cumulative distribution function for incomes of the entire population, $g_J(x)$ is the density function for incomes of members of the group, *J*, and the integrals are calculated over the domain of all possible incomes.

Appendix 2 (Proof of Lemma 1)

Without loss of generality, we shall assume that the attributes can be understood as income levels. Let $\boldsymbol{A}$ refer to a matrix of size *K* by *n* with *K* identity groups and *n* income levels. Each element in the matrix $a_{ij} \in \{0, 1, 2\}, \forall (i, j)$. We say that the *ith* identity group is “under-represented” at the *jth* income level if the proportion of persons from group *i* at income *j* is less than the proportion of persons of group *i* in the population as a whole. We denote the statement that the *ith* identity group is “under-represented” at the *jth* income level by $a_{ij} = 0$. We say further that the *ith* identity group is “over-represented” if the proportion of persons from group *i* at income *j* is greater than the proportion of persons of group *i* in the population as a whole. We denote the statement that the *ith* identity group is “over-represented” at the *jth* income level by $a_{ij} = 1$. If the *ith* identity group is represented at the *jth* income level in the same proportion as it is represented in the population as a whole, then we say that it is "equiproportionally represented" and we denote this by $a_{ij} = 2$.

Thus, $\boldsymbol{A}$ is a matrix in which every element is 0,1, or 2. We may further note that if any row or any column contains a zero, then it must contain a one and vice versa. This requirement captures the necessity that if an identity group is over-represented at an income level, it must be under-represented at another income level and that if a group is over-represented at an income level, then another group is under-represented at that same income level.

A balanced bilateral transfer is always possible if an identity group is represented to a greater extent at one income level (call it y_1) than it is at another (call it y_2) and another identity group is represented to a lesser extent at y_1 than it is at y_2. This condition is satisfied as long as it is possible to identify two rows (i and j) and two columns (l and m) of the matrix A such that they form a matrix $A^{\sim} = \begin{pmatrix} a_{il} & a_{im} \\ a_{jl} & a_{jm} \end{pmatrix}$ which is of one of the following forms, or which can be constructed from one of the following forms by permuting either their rows or their columns:

$$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}, \begin{pmatrix} 1 & 2 \\ 2 & 1 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 2 \\ 1 & 0 \end{pmatrix}.$$

The lemma is therefore equivalent to the statement that there exists a matrix $\boldsymbol{A}^{\sim}$ for any matrix $\boldsymbol{A}$ which contains at least a single one or zero. Suppose that the lemma is false. Then, it is possible to construct an $\boldsymbol{A}$ such that there is no $\boldsymbol{A}^{\sim}$ associated with it.

We now try to construct such a matrix $\boldsymbol{A}$. Without loss of generality, consider the case in which **A** contains at least one zero (i.e., an identity group is under-represented at a particular level of income). We can present this as occurring at the top left corner (a_{11}) of the matrix, without loss of generality, as given below:

$$\boldsymbol{A} = \begin{pmatrix} 0 & \dots & a_{1n} \\ \vdots & \ddots & \vdots \\ a_{k1} & \cdots & a_{kn} \end{pmatrix}.$$

This however means that there must be at least one level of income in the first column and in the first row in which there is over-representation of an identity group. Without loss of generality, let us say that this occurs at a_{12} and a_{21}, respectively, so that

$$\boldsymbol{A} = \begin{pmatrix} 0 & 1 & \dots a_{1n} \\ 1 & a_{22} & \cdots a_{2n} \\ \vdots & \dots & \ddots \\ & & ak1 \end{pmatrix}.$$

Now, if $a_{22} = 0$ or 2, then $\boldsymbol{A}^{\sim}$ exists. If $a_{22} \neq 0$ or 2, then $a_{22} = 1$. That is

$$\boldsymbol{A} = \begin{pmatrix} 0 & 1 & \ldots a_{1n} \\ 1 & 1 & \cdots a_{2n} \\ \vdots \ldots & & \ddots \\ & & ak1 \end{pmatrix}.$$

Consider row 2 and column 2 now. Since for the already fixed elements, there is over-representation, there must be elements in row 2 and in column 2, respectively, that have value zero (reflecting under-representation). Without loss of generality, let these occur at a_{23} and a_{32}, respectively, so that

$$\boldsymbol{A} = \begin{pmatrix} 0 & 1 & a_{13}\ldots \\ 1 & 1 & 0\ldots\ldots \\ a31 & 0 & a33\ldots \\ & & \vdots \end{pmatrix} = \begin{pmatrix} 0\,1\,0\ldots\ldots \\ 1\,1\,0\ldots\ldots \\ 0\,0\,0\ldots\ldots \\ \vdots \end{pmatrix}.$$

But this in turn fixes a_{13}, a_{31}, and a_{33} to be 0 since if any of these are 1 or 2, we can construct matrix $\boldsymbol{A}^{\sim}$. This in turn implies that there exist elements elsewhere in row 3 and column 3 with value 1 (indicating over-representation), which we can place without loss of generality at a_{34} and a_{43}, respectively. It can readily be seen that this in turn fixes a_{41}, a_{42}, a_{44}, a_{24}, and a_{14} to be 1 since if any of these are 0 or 2, we can construct matrix $\boldsymbol{A}^{\sim}$. Thus we may construct a matrix $\boldsymbol{A}$ such that $a_{ij} = a_{ji}$ $= 0$, if i is odd and $j \leq i$ and $a_{ij} = 1$ otherwise.

Let us now consider the matrix where the row $(k - 1)$ is odd. This means that $\boldsymbol{A}$ has the following form:

$$\boldsymbol{A} = \left(\begin{array}{ccccc|c} \vdots & \vdots & \vdots & \vdots & \vdots & a1n \\ \cdots 1 & 1 & 0 & 1 & 0 & \vdots \\ \cdots 0 & 0 & 0 & 1 & 0 & \vdots \\ \cdots 1 & 1 & 1 & 1 & 0 & \vdots \\ \cdots 0 & 0 & 0 & 0 & 0 & \vdots \\ ak1 & \cdots & \cdots & \cdots & \cdots & akn \end{array}\right)$$

This in turn implies that $a_{k-1,n} = 1$ and $a_{k,n-1} = 1$. It may be verified that for $\boldsymbol{A}^{\sim}$ not to exist all elements in row k and in column n must equal 1. However, this violates the requirements on a matrix $\boldsymbol{A}$.

Consider now the matrix where the row $(k - 1)$ is even. This means that $\boldsymbol{A}$ has the following form:

$$
\boldsymbol{A} = \left(\begin{array}{ccccc|c}
\vdots & \vdots & \vdots & \vdots & \vdots & a1n \\
\cdots 0 & 0 & 1 & 0 & 1 & \vdots \\
\cdots 1 & 1 & 1 & 0 & 1 & \vdots \\
\cdots 0 & 0 & 0 & 0 & 1 & \vdots \\
\cdots 1 & 1 & 1 & 1 & 1 & \vdots \\
ak1 & \cdots & \cdots & \cdots & \cdots & akn
\end{array}\right)
$$

This in turn implies that $a_{\mathrm{k}-1,\mathrm{n}} = 0$ and $a_{\mathrm{k},\mathrm{n}-1} = 0$. It may be verified that for $\boldsymbol{A}^{\sim}$ not to exist all elements in row k and in column n must equal 0. However, this violates the requirements on a matrix $\boldsymbol{A}$.

Thus, it is not possible to construct a matrix $\boldsymbol{A}$ such that $\boldsymbol{A}^{\sim}$ does not exist.

$A^{\sim}$ must exist, thereby proving the lemma.

QED

References

Abul Naga RH, Yalcin T (2009) Inequality Measurement for Ordered Response Health Data. J Health Econ 27(6):1614–1625

Alesina A, La Ferrara E (2000) Participation in heterogeneous communities. Q J Econ 115, 847–904

Alesina A, La Ferrara E (2002) Who trusts others? J Public Econ 85, 207–34

Alesina A, Devleeschauwer A, Easterly W, Kurlat S, Wacziarg R (2003) Fractionalization. J Econ Growth 8, 155–194

Allison RA, Foster J (2004) Measuring health inequalities using qualitative data. J Health Econ 23:505–524

Anderson E (1999) What is the point of equality? Ethics 1999:287–337

Anderson G (2004) Toward an Empirical Analysis of Polarization. Journal of Econometrics 122(2004):1–26

Anderson G (2005) Polarization. Working Paper, University of Toronto. Available at http://www.chass.utoronto.ca/~anderson/Polarization.pdf

Arneson R (1989) Equality and equal opportunity for welfare. Philos Stud 1989:77–93

Bartolucci F, Forcina A, Dardanoni V (2001) Positive quadrant dependence and marginal modelling in two-way tables with ordered margins. J Am Stat Assoc 96(456), 1497–1505

Boisso D, Hayes K, Hirschberg J, Silber J (1994) Occupational segregation in the multidimensional case: decomposition and tests of statistical significance. J Econ 61:161–171

Chakravarty S, Maharaj B (2009) A study on the RQ index of ethnic polarization. ECINEQ WP 134

Cohen GA (1989) On the currency of egalitarian justice. Ethics 1989:906–944

Duclos J-Y, Esteban J, Ray D (2004) Polarization: concepts, measurement, estimation. Econometrica 1737–1772

Duncan OD, Duncan B (1955) A methodological analysis of segregation indexes. Am Sociol Rev 20(2):210–217
Dworkin R (2000) Sovereign virtue: the theory and practice of equality. Harvard University Press, Cambridge, MA
Esteban J-M, Ray D (1994) On the measurement of polarization. Econometrica 62:819–851
Foster J, Shneyerov A (1999) A general class of additively decomposable inequality measures. Econ Theory 14(1)
Flückiger Y, Silber J (1994) The gini index and the measurement of multidimensional inequality. Oxf Bull Econ Stat 56:225–228
Hutchens R (1991) Segregation curves, lorenz curves, and inequality in the distribution of people across occupations. Math Soc Sci 21:31–51
Hutchens R (2004) One measure of segregation. Int Econ Rev 45(2):555–578
Jayaraj D, Subramanian S (2006) Horizontal and vertical inequality: some interconnections and indicators. Soc Indic Res 75(1):123–139
Jayadev A, Reddy S (2011) Inequalities between groups: theory and empirics. World Dev 39(2), 159–173
Loury GC (2004) The anatomy of racial inequality. Harvard University Press, Cambridge, MA
Miguel E, Gugerty M (2005) Ethnic diversity, social sanctions, and public goods in Kenya. J Public Econ 89(11–12):2325–2368
Montalvo JG, Reynal-Querol M (2005) Ethnic polarization, potential conflict and civil war. Am Econ Rev 95 (3), 796–816
Østby Gudrun (2008) Polarization, horizontal inequalities and violent civil conflict. J Peace Res 45(2):143–162
Putnam H (2004) The collapse of the fact/value dichotomy and other essays. Harvard University Press, Cambridge MA
Rawls J (1971) A theory of justice. Harvard University Press, Cambridge, MA
Reardon SF, Firebaugh G (2002) Measures of multigroup segregation. Sociol Methodol 32:33–67
Reardon S (2008) Measures of ordinal segregation. Working paper 2008–2011, Institute for Research on Education Policy and Practice, Stanford University. Forthcoming in Research on Economic Inequality, vol 17
Reddy S, Jayadev A (2011) Inequalities and identities. Available at SSRN: https://ssrn.com/abstract=1162275 or http://dx.doi.org/10.2139/ssrn.1162275
Roemer J (1996) Theories of distributive justice. Harvard University Press, Cambridge, MA
Schweizer B, Wolff EF (1981) On nonparametric measures of dependence for random variables. Ann Stat 9(4):879–885
Sen A (2007) We can best stop terror by civil, not military, means. The Guardian, Friday November 9th, 2007
Sen A (1992) Inequality reexamined. Harvard University Press, Cambridge, MA
Shorrocks AF (1980) The class of additively decomposable inequality measures. Econometrica 48(3):613–625
Shorrocks AF (1983) Ranking income distributions. Economica 50:3–17
Silber JG (1989) On the measurement of employment segregation. Econ Lett 30(3):237–243. Elsevier
Silber JG (1991) Inequality indices as measures of dissimilarity: a generalization. Stat Pap 1991(32):223–231
Silber JG (1992) Occupational segregation indices in the multidimensional case: a note. Econ Rec 68:276–277
Stewart F (2001) Horizontal inequalities: a neglected dimension of development. WIDER Annual Lectures 5, UNU WIDER
Tchen AH (1980) Inequalities for distributions with given marginals. Ann Probab 8(4):814–827
Wolfson MC (1994) When inequalities diverge. Am Econ Rev Pap Proc 94:353–358
Zhang X, Kanbur R (2001) What difference do polarisation measures make? an application to china. J Dev Stud 37:85–98