J Econ Inequal (2008) 6:231–245
DOI 10.1007/s10888-007-9064-x

Reinterpreting between-group inequality

Chris Elbers · Peter Lanjouw · Johan A. Mistiaen · Berk Özler

Received: 7 April 2006 / Accepted: 4 May 2007 /
Published online: 1 September 2007

Abstract We evaluate observed inequality between population groups against a benchmark of the maximum between-group inequality attainable given the number and relative sizes of those groups under examination. Because our measure is normalized by these parameters, drawing comparisons across different settings is less problematic than with conventional inequality decompositions. Moreover, our measure can decline with finer sub-partitioning of population groups. Consequently, the exact manner in which one groups the population acquires greater significance. Survey data from various countries suggest that our approach can provide a complementary perspective on the question of whether (and how much) a particular population breakdown is salient to an assessment of inequality in a country.

Keywords Inequality decompositions

JEL Classification D31

1 Introduction

The significance of group differences in wellbeing is often at the center of the study of inequality. Statistical methods that are often used to 'decompose' economic inequality into

C. Elbers (✉)
Free University, Amsterdam, The Netherlands
e-mail: celbers@feweb.vu.nl

C. Elbers
Tinbergen Institute, Amsterdam, The Netherlands

P. Lanjouw · J. A. Mistiaen · B. Özler
World Bank, Washington, DC, USA

P. Lanjouw
e-mail: planjouw@worldbank.org

J. A. Mistiaen
e-mail: jmistiaen@worldbank.org

B. Özler
e-mail: bozler@worldbank.org

constituent parts are well-known to economists. Sub-group decomposable measures of inequality can be written as the sum of inequality that is attributable to differences in mean outcomes across population sub-groups and that which is due to inequality *within* those sub-groups.[1] Many have used such decompositions to 'understand' economic inequality and guide the design of economic policy. Indeed, Cowell [4] argues: "It is almost essential to attempt to 'account for' the level of, or trend in, inequality by components of the population."

Conventionally, between-group inequality depends on three factors: differences among groups in mean incomes, the number of the groups, and their relative sizes. Because underlying population structures often vary, this causes difficulties with comparisons of such decompositions across different settings. Consider three countries where the issue of racial differences in income features prominently in public discourse: the United States, Brazil and South Africa. The shares of income inequality attributable to differences between racial groups in these countries are 8%, 16%, and 33%, respectively.[2] Do these numbers provide a good yardstick with which to judge the relevance of race to an understanding of inequality in these countries? Should South African and Brazilian policy-makers worry much more about racial differences in incomes than do their American counterparts? Does the small percentage of income inequality attributable to race in the U.S. mean that racial inequality is not a pertinent economic and social issue?

The figures above are based on four population groups for Brazil and South Africa, and five for the U.S., but the population shares of the white groups versus non-white groups differ tremendously.[3] In each country, the mean income of the non-white groups is much below that of the white group, but the non-white groups form the majority in South Africa (80%), half of the population in Brazil (50%), and a minority in the U.S. (28%). The difference in between-group inequality observed between these three countries could in fact be due largely to the difference in population shares of the racial groups instead of the differences in relative mean incomes of these groups.[4] Hence, the first difficulty with inequality decompositions is caused by the fact that they are not really comparable because they are not unit-free: they depend on the number and relative sizes of the groups under examination.

A second issue concerns the interpretation of inequality decompositions and their implications for policy design. Although decompositions of inequality have long been the workhorse in this literature, empirical implementation has tended to find little evidence of significant between group differences. For example, in a classic reference, Anand [1]

[1] See Bourguignon [2], Shorrocks [13, 14] and Cowell [3]. Cowell [4] provides a recent survey of methods of inequality measurement, including a discussion of the various approaches to sub-group decomposition. In this paper we focus on the class of inequality measures that can be additively decomposed into a within-group and a between-group component. As the Gini coefficient does not lend itself to such a neat decomposition, we will not be focusing attention on this measure. Elbers et al. [8] provides some discussion of how the findings of this paper bear on decompositions of the Gini coefficient.

[2] These figures have been calculated by the authors using data from PNAD (2001) for Brazil, IES(2000) for South Africa, and LIS(2000) for the U.S.

[3] The racial groups used in our analysis are "White", "Black", "Pardo" and "indigenous" in Brazil, "African", "Colored", "Asian/Indian" and "White" in South Africa, and "White", "Black", "Hispanic", "Asian", and "American Indian" in the U.S.

[4] The observed differences in between-group inequality may also depend on the number of groups under consideration, making the specific definition of groups a non-negligible issue. For example, the share of between-group inequality attributable to caste in India when one groups people simply into "high", 'medium", or "low" caste groupings, could be quite different from that which emerges when the partitions are *finer*, i.e. when one makes distinctions between castes within each broad category.

showed that inequality between ethnic groups in Malaysia accounted for only 15% of total inequality in the 1970s. This led to his recommendation that government strategy should focus on inequality within ethnic groups rather than that between them. Cowell and Jenkins [5], who find that most income inequality remains unexplained even after taking into account the age, sex, race and earner status of the household head in the U.S., argue that the real story of inequality is to be found *within* these population groups and point to the importance of chance.[5]

Not everyone is comfortable with such interpretations, however. Kanbur [10] states that the use of such decompositions "...assists the easy slide into a neglect of inter-group inequality in the current literature." He argues that finding a relatively small share of inequality between groups does not mean that the mean differences between them are less important than inequalities within such groupings. In particular, he argues that social stability and racial harmony can break down once the average differences between groups go beyond a certain threshold, with the threshold varying from country to country.[6]

Perhaps, it is not so surprising that one rarely observes a high share of between-group inequality. The conventional between-group share is calculated by taking the ratio of observed between-group inequality to total inequality. Total inequality, however, can be viewed as the between-group inequality that would be observed if every household in the population constituted a separate group. Thus, the conventional practice is equivalent to comparing observed between-group inequality (across a few groups under examination) against a benchmark (across perhaps millions of groups) that is quite extreme—and probably rather unrealistic.

In this paper, we address these two difficulties in interpreting inequality between groups, namely *comparability* and the rather *extreme benchmark* against which between-group inequalities are judged, by proposing an alternative measure. Specifically, we suggest replacing total inequality in the denominator of the conventional ratio with the maximum between-group inequality that could be obtained if the number of groups and their sizes were restricted to be the same as for the numerator. Because our proposed measure is normalized by the number of groups under examination and their relative sizes, one can more readily make comparisons across settings where the number of groups is (or the population shares for those groups are) very different.

The paper is organized as follows. Section 2 defines our new measure and discusses available estimation and computation approaches. Section 3 uses household survey data from various countries to provide assessments of between-group inequality based on the conventional method and contrasts the conclusions one might draw with those when the analysis is based on our alternative measure. Qualitative assessments of the importance of

---

[5] Elbers et al. [7] find that the share of consumption inequality between the 915 rural sub-locations of Ecuador is only 14% and that between 1,117 rural sub-locations in Madagascar is 18%. As discussed above, after almost half a century of racial segregation and oppression, inequality between races still accounts for only a third of total consumption inequality in South Africa.

[6] "Sub-group consistency requires that a change in a subgroup's distribution which happens to raise inequality in the subgroup must lead to an overall increase in inequality, no matter how that change influences the relative positions of the remaining population." (Foster and Sen [9], page 159) Foster and Sen point out that this 'rather separatist' view implicit in these sub-group consistent measures ignores potentially relevant information when making inequality comparisons. For example, should a change in inequality within a certain group (while the means and population shares remain unchanged) when that group is richer than a second group affect inequality in exactly the same manner as in the presence of a much richer second group? Sub-group consistency requires this to be true. Kanbur [10] builds on this argument and suggests that invoking such axioms "...go[es] against basic intuition and considerable evidence which suggest that individuals do indeed pay special attention to outcomes for their particular racial, ethnic, or regional group."

between-group differences can indeed be markedly different when based on this alternative approach. This section also discusses a thought-provoking finding of a strong, positive cross-country correlation between overall inequality and between-group inequality. Section 4 concludes.

2 Methodology

Given a partition of the population Π, additively decomposable inequality measures can be written as follows:[7]

$$I = I_w\left(\prod\right) + I_B\left(\prod\right)$$

where $I_w(\prod)$is a weighted average of inequality within population sub-groups, while $I_B(\prod)$\stands for between-group inequality and can be interpreted as the amount of inequality that would be found in the population if everyone were given the average income of their group.

The most commonly decomposed measures in this literature come from the General Entropy class. These take the following form:

$$\begin{aligned} GE &= \frac{1}{c(c-1)} \sum_i f_i \left[\left(\frac{y_i}{\mu} \right)^c - 1 \right] && \text{for } c \neq 0, 1 \\ &= \sum_i f_i \log \left(\frac{\mu}{y_i} \right) && \text{for } c = 0 \\ &= \sum_i f_i \left(\frac{y_i}{\mu} \right) \log \left(\frac{y_i}{\mu} \right) && \text{for } c = 1 \end{aligned}$$

where f_i is the population share of household i, y_i is per capita consumption of household i, μ is average per capita consumption, and c is a parameter that is to be selected by the user.[8] This class of inequality measures can be neatly decomposed into a between- and within-group component as follows [6, 11, 14]:

$$\begin{aligned} GE &= \frac{1}{c(c-1)} \left[\sum_j g_j \left(\frac{\mu_j}{\mu} \right)^c - 1 \right] + \sum_j GE_j g_j \left(\frac{\mu_j}{\mu} \right)^c && \text{for } c \neq 0, 1 \\ &= \left[\sum_j g_j \log \left(\frac{\mu}{\mu_j} \right) \right] + \sum_j GE_j g_j && \text{for } c = 0 \\ &= \left[\sum_j g_j \left(\frac{\mu_j}{\mu} \right) \log \left(\frac{\mu_j}{\mu} \right) \right] + \sum_j GE_j g_j \left(\frac{\mu_j}{\mu} \right) && \text{for } c = 1 \end{aligned}$$

where j refers to the sub-group, g_j refers to the population share of sub-group j and GE_j refers to inequality in sub-group j. The between-group, $I_B(\Pi)$, component of inequality is captured by the first term: the level of inequality if everyone within each sub-group j had consumption level μ_j. The second term gives within-group inequality $I_w(\prod)$.

[7] We borrow our notation mainly from Cowell and Jenkins [5].

[8] Lower values of c are associated with greater sensitivity to inequality amongst the poor, and higher values of c place more weight to inequality among the rich. A c value of 1 yields the well known Theil entropy measure, a value of 0 provides the Theil L or mean log deviation, and a value of 2 is ordinally equivalent to the squared coefficient of variation.

Given a partition Π and an inequality measure I, between-group inequality can be summarized as follows:

$$R_B\left(\prod\right) = \frac{I_B(\prod)}{I}$$

$R_B(\prod)$ represents the share of inequality explained by between-group differences. For any characteristics x and y, $R_B(\prod_{x\ \&\ y}) \geq R_B(\prod_x)$ and $R_B(\prod_{x\ \&\ y}) \geq R_B(\prod_x)$.[9] This means that moving from any partition to a finer sub-partition, the share of between-group inequality cannot decrease.

2.1 Maximum between-group inequality

Using the notion of between-group inequality described above, it is not uncommon to encounter statements of the following type: "inequality between groups accounts for *only* 20% of the total inequality in incomes." Such statements, however, should not be taken to mean that 100% of total inequality would have been a realistic possibility. Between-group inequality would equal total inequality under only two unlikely scenarios: (i) if each household itself constituted a group, or (ii) if there were fewer groups than households, but somehow all the households within each of these groups happened to have identical per capita incomes. It is difficult to imagine a realistic setting in which either of these scenarios would occur: for virtually any empirically relevant income distribution and a limited number of groups (much smaller than the number of individuals in the population), the share of maximum between-group inequality that can be attained is strictly below unity.

While assessing the importance of between-group inequality for a certain partition, if total inequality is an extreme benchmark then what is a relevant one? In this paper, we propose that one possibility is to evaluate observed between-group inequality for a certain partition against a benchmark of maximum between-group inequality that can be attained when the number and relative sizes of groups for that partition are unchanged. In other words, we propose to compare actual observed between-group inequality against a counterfactual between-group inequality constructed from the same data, using the same number of groups and relative sizes, but where households in the income distribution are re-assigned to the population groups in such a manner so as to maximize between-group inequality.

The index we propose is defined as:

$$\widehat{R}_B(\prod) = \frac{I_B(\prod)}{Max\{I_B|\prod(j(n), J)\}} = R_B\left(\prod\right)\frac{I}{Max\{I_B|\prod(j(n), J)\}},$$

where the denominator is the maximum between-group inequality that could be obtained by reassigning individuals across the J sub-groups in partition Π of size $j(n)$.

Since between-group inequality can never exceed total inequality, it follows that $\widehat{R}_B(\prod)$ cannot be smaller than $R_B(\prod)$. However, unlike the traditional between-group inequality measure, our alternative measure, $\widehat{R}_B(\prod)$, does not necessarily increase when a finer partitioning is obtained from the original one. This is because, for $\widehat{R}_B(\prod)$, both the numerator and the denominator change as a result of finer partitioning, and whether it

[9] See, for example, Cowell and Jenkins [5].

increases or not depends on the relative rate of change of these two components.[10] Furthermore, for any finer partitioning of an original partition, the (signed) rate of change of $\widehat{R}_B(\prod)$ is lower than (or equal to) that of $R_B(\prod)$. The proof is obvious: the rate of change in the numerator is the same for both measures, but while the rate of change in the denominator for $R_B(\prod)$ is zero, it is nonnegative for $\widehat{R}_B(\prod)$. This implies that $\widehat{R}_B(\prod)$ cannot proportionally increase faster than $R_B(\prod)$ for any finer partitioning—in fact it may decline. The two properties of our measure described above imply (i) that there may be a large difference between $R_B(\prod)$ and $\widehat{R}_B(\prod)$ for a particular partition Π, but also (ii) that they must converge towards each other with each successive sub-partitioning.

2.2 Calculating $\widehat{R}_B(\prod)$

In order to calculate $\widehat{R}_B(\prod)$ we need to know $I_B(\prod)$, which can be calculated in the usual way, and maximum between-group inequality, which is slightly more difficult to compute. Calculating maximum between-group inequality uses the property that under a between-group inequality maximizing distribution, sub-group incomes occupy non-overlapping intervals. This is a necessary condition for between-group inequality to be at its maximum: if $\{y\}$ is an income distribution for which inequality between sub-groups g and h is maximized, then either *all* incomes in are higher than *all* incomes in h, or *vice versa* (see Shorrocks and Wan [15], section 3).

In the case of J sub-groups in a particular partition, in principle the following approach can be followed: take a particular permutation of sub-groups $\{g_{(1)},..., g_{(J)}\}$, allocate the lowest incomes to $g_{(1)}$, then to $g_{(2)}$, etc., and calculate the corresponding between-group inequality. Repeat this for all possible $J!$ permutations of sub-groups.[11] The highest resulting between-group inequality is the maximum sought.

A possibly more appealing benchmark against which to evaluate between-group inequality can be obtained by introducing one more restriction. In addition to fixing the number of sub-groups and their relative sizes, we can also arrange the sub-groups under examination according to their observed mean incomes, keeping their 'pecking order' unchanged.[12] In many cases, there is a well-understood hierarchy of population groups in terms of their mean incomes. Comparing actual between-group inequality to a counterfactual maximum that preserves the actual, observed, rank ordering of sub-groups is conceivably of greater interest than a counterfactual that allows for random re-orderings of the sub-groups. For example, when decomposing inequality by race in Brazil, South Africa, or the U.S. (see the example in Sect. 1), the ordering of racial sub-groups in terms of mean incomes is well-documented, and it is not obvious to what extent a counterfactual of say, average income of blacks exceeding that of whites would be realistic and of any inherent interest.

Obtaining the maximum possible between-group inequality given the current income distribution, relative sub-group sizes, *and* their rankings by mean incomes is also simpler because we need to calculate between-group inequality only once instead of $J!$ times for all

[10] It is relatively easy to construct examples where $\widehat{R}_B$ increases or decreases with finer partitioning. See Section 3 for such examples.

[11] Obviously, this approach requires the number of groups, J, to be relatively small.

[12] Ordering population groups by their mean incomes using, say, a household survey would introduce a possible difficulty due to sampling variability. In other words, our ability to order groups by mean income (or consumption) could be limited by the fact that some of the group means are statistically indistinguishable from each other. For the time being, we ignore the standard errors associated with the observed group means.

possible orderings of the sub-groups. An example might clarify: suppose we have three sub-groups with population shares of 50%, 30%, and 20%, respectively. The largest sub-group has the lowest observed mean income, and the smallest the highest. Maximum between-group inequality (given 'pecking order') is obtained by generating three sub-groups with non-overlapping incomes, where the poorest sub-group occupies the bottom half of the distribution, the next sub-group occupies the incomes between the median and the 80th percentile, and the final sub-group the top 20% of the income distribution. In the empirical section of this paper that follows, $\widehat{R}_B(\prod)$ will refer to our index of between-group inequality normalized by the maximum possible between-group inequality given the current income distribution, relative sub-group sizes, *and* their 'pecking order.'[13]

It is important to note that $\widehat{R}_B(\prod)$ is not the product of a strict decomposition exercise.[14] As such, we view $\widehat{R}_B(\prod)$ not as an alternative to $R_B(\prod)$, but as a complement. Knowledge of $\widehat{R}_B(\prod)$ can aid in the assessment and interpretation of the importance of inequality between sub-groups in various settings. For example, Cowell and Jenkins [5] argue that "if, for two alternative partitions Π_x and Π_y corresponding to two population characteristics x and y, we find that $R_B(\prod_x)$ is much greater than $R_B(\prod_y)$, then it is evidently reasonable to say that in some sense the population characteristic x is more important as a determinant of inequality than is characteristic y." Given difficulties of *comparability*, however, such a conclusion may not be so evident. As we will see in the next section, it is possible for $R_B(\prod_x)$ to be greater than $R_B(\prod_y)$, *and* for $\widehat{R}_B\left(\prod_y\right)$ to be greater than $\widehat{R}_B\left(\prod_x\right)$. In such circumstances, the interpretation of the importance of population characteristics x and y might be different than when that assessment is based only on the conventional decomposition methods.

3 Evidence

Using household survey data from eight countries, Table 1 presents total inequality in consumption expenditures, the conventional share of between-group inequality, and our proposed measure, where the sub-groups are defined by the relevant racial, ethnic, or caste breakdown in each country.[15] For example, the breakdown for the United States corresponds to five racial sub-groups: Whites, Blacks, American Indians, Asians and Hispanics. In India, the three sub-groups comprise Scheduled Caste households, Scheduled Tribes, and Others. The number of sub-groups and their respective sizes are clearly not the same in all countries. Inequality is measured on the basis of per-capita consumption for each country and we have chosen the General Entropy Class measure with parameter value zero, also referred to as the Theil L measure or the mean log deviation.

Based on the standard approach to decomposing inequality, as described above, between-group inequality in each country in our list is rather low. Only South Africa stands out with a conventional between-race share (R_B) of 33%, although even here it is striking to

[13] The maximum between-group inequality possible when the 'pecking order' of groups is kept fixed will always be less than (or equal to) that over $J!$ permutations. Consequently, the value of $\widehat{R}_B(\prod)$ can be different under these two methods.

[14] Obviously, our proposed measure of between-group inequality and within-group inequality do not add up to total inequality.

[15] The eight countries were selected on the basis of availability of household surveys with consumption expenditure data and suitable identification of social groups.

Table 1 Decomposing inequality by "Social" Group in 8 countries

Country	No. of "social" groups	GE(0)	R_B	$\widehat{R}_B$
India	3	0.136	5.1	10.1
Bangladesh	4	0.181	20.3	28.7
Kazakhstan	3	0.217	9.0	14.7
Nepal	10	0.220	23.3	23.7
United States	5	0.295	8.4	14.7
Panama	7	0.402	13.8	31.8
Brazil	4	0.408	15.8	21.6
South Africa	4	0.607	33.3	56.4

Data for India refer to rural areas only. Social group refers to the relevant racial, ethnic, or caste breakdown in each country. GE(0) refers to the mean log deviation in per-capita consumption for each country. R_B is the conventional share of between-group inequality in total inequality, while $\widehat{R}_B$ is our proposed measure.

note that two-thirds of total inequality in a country that suffered nearly half a century of racial segregation can be attributed to differences *within* racial sub-groups as opposed to differences *across* them. However, using our alternative measure $\left(\widehat{R}_B\right)$, we find that observed inequality between the four racial sub-groups accounts for more than 56% of the 'maximum possible' between-race inequality in South Africa given its current income distribution, the number of racial sub-groups, their sizes, and their ranking in terms of average income. As detailed in the previous section, our measure would take the value 0 if all group means were identical, and 1 if none of the group distributions overlapped with each other. Hence, in South Africa, the current distribution of income between racial groups is more than halfway towards a completely segregated distribution of incomes on this spectrum.

A slightly different observation can be made by examining the figures for Brazil and Panama in Table 1. Based on the standard decomposition by race/ethnicity, the between-group share of inequality in both countries is less than 16%.[16] This can conventionally be interpreted as suggesting that race or ethnicity is of limited relevance to an understanding of inequality in these two countries.[17] However, in Panama, observed inequality between ethnic sub-groups accounts for about a third of 'maximum possible' inequality between such sub-groups, while in Brazil the conclusion based on our measure is only slightly different from that which is obtained from the standard calculation.

Comparisons of the conventional and alternative approaches can also be instructive when examining the importance of different characteristics within the same country. Table 2 presents the conventional share of between-group inequality, and our proposed measure, for three different household characteristics in Thailand for 2002: whether the household lives in an urban or rural area, its geographic region, and the education level of the head of the household. It also presents these measures for all sub-partitions that can be formed by combining these sub-groups. For example, we can examine inequality between sub-groups

[16] The social groups in Panama are based on the language spoken at home. Spanish speakers constitute 90% of the population.

[17] Of course for a characteristic, such as race, ethnicity, or gender, that is a circumstance rather than one that is related at least in part to individual choice, such as level of education or occupation, any positive between-group inequality could be viewed as too high and unacceptable.

Table 2 Decomposing inequality in Thailand

Grouping	No. of "social" groups	Percentage share of the largest group	R_B	$\widehat{R}_B$
Urban/rural (U)	2	69.7	23.5	36.1
Education level (E)	4	70.9	29.4	36.6
Region (R)	5	34.3	25.9	28.2
U×E	8	54.8	39.7	43.0
U×R	9	29.0	33.5	34.4
R×E	20	28.1	44.5	45.5
U×R×E	36	24.9	47.5	47.8

R_B is the conventional share of between-group inequality in total inequality (measured by mean log deviation in per-capita consumption), while $\widehat{R}_B$ is our proposed measure. Thailand has 5 regions, but because Bangkok has no rural areas the breakdown of regions into urban and rural areas yields only nine groups instead of 10. We created four groups for the education level of the household head: none, primary, secondary, and post-secondary.

where the sub-groups are defined by the four levels of education in each of the five regions, yielding 20 sub-groups.

The first thing to notice in Table 2 is the rank reversals of population characteristics when we switch from the conventional decomposition to our proposed measure. R_B is 29.4% for education level and 23.5% for urban/rural, while $\widehat{R}_B$ is about 36% for both of them. Consumption inequality between regions is 25.9% when measured by R_B, but only 28.2% when measured by $\widehat{R}_B$—significantly lower than the 36.1% between urban and rural areas. It does seem that living in an urban or rural area and education level of the head of the household are more salient characteristics correlated with inequality than geographic location in Thailand.

The second thing to note in Table 2 is related to a major difference in how R_B and $\widehat{R}_B$ behave. As we can see, with finer partitioning, R_B monotonically increases regardless of the order of characteristics by which we decompose inequality. However, as mentioned in the previous section, $\widehat{R}_B$ can decline with finer partitioning. We find that starting from an urban/rural grouping only (36.1%) to a combination of urban/rural and region (34.4%) to the combination of all three characteristics (47.8%), $\widehat{R}_B$ declines before increasing again. Figures 1 and 2 demonstrate this difference in the properties of the two measures more clearly.[18] Depending on the order in which the population is partitioned, the two measures can chart different paths. Regardless, however, $\widehat{R}_B$ always starts above R_B and necessarily follows a flatter slope until the two measures converge (when the partitioning is sufficiently fine). Note that once the largest sub-group is sufficiently small (in this case accounting for less than 30% of the population), R_B and $\widehat{R}_B$ are roughly identical. The differences between the two measures can be significant when one of the sub-groups accounts for a large share of the population, which is usually the case when the number of sub-groups in a partition is small.

The fact that $\widehat{R}_B$ can decline for finer sub-partitioning of a certain population sub-group points us in an interesting yet relatively uncharted direction. Many researchers who make use of the conventional inequality decomposition tools are preoccupied with the question of

[18] In the figures, R_B and $\widehat{R}_B$ are referred to as *conventional* and *alternative* between-group shares, respectively.

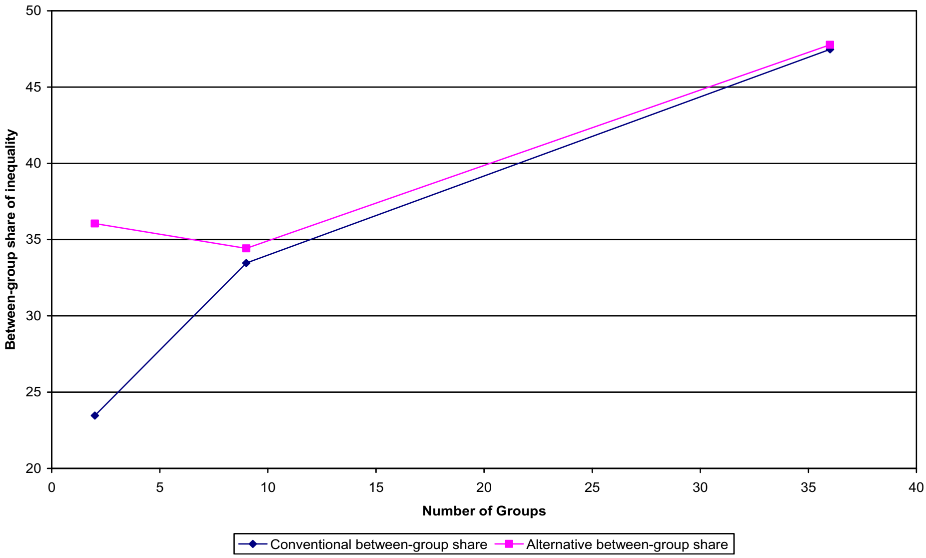

Fig. 1 Finer partitions: urban/rural, then region, then education

how much of the existing inequality we can explain by grouping the population in various ways.[19] The idea is that the finer the grouping, the more inequality we can attribute to differences between the sub-groups.[20] However, as demonstrated above, this is clearly not the case when we use $\widehat{R}_B$ to examine inequality between sub-groups. It is possible that finer partitioning of the population for a certain characteristic does not add much to the picture, but may in fact dilute it. This would be the case if the sub-groups created under the sub-partition are not as 'salient' to inequality as the original, coarser, sub-groups.

Take inequality between ethnic groups in South Africa. One could look at inequality between simply 'Blacks' and 'Whites'.[21] We could also examine inequality between racial groups for another partitioning of the population: Africans and non-Africans. Africans are the largest (accounting for roughly 80% of the population) and the poorest of the four racial groups in South Africa [12]. We could investigate inequality between the four major racial groups, making a distinction between Africans, Coloreds, and Asians. Finally, we could further make a distinction on language spoken at home, which for the African population is an indicator of the particular ethnic group the household belongs to. Table 3 presents R_B and $\widehat{R}_B$ for these alternative partitions of the population by racial group.

Examining the conventional inequality decomposition measure, R_B, we find that it is about 27% whether we break the population into "Whites/non-Whites" or "Africans/Non-Africans". However, the measure we propose, $\widehat{R}_B$, is 80% for the former breakdown while only 50% for the latter. What explains the very high inequality (measured by $\widehat{R}_B$) between

[19] In fact, this is the title of the paper by Cowell and Jenkins [5]: "How Much Inequality Can we Explain? A Methodology and an Application to the United States".

[20] For example, Cowell and Jenkins [5] combine four household characteristics to examine inequality between 128 groups. Elbers et al. [7] examine inequality between hundreds of small communities in three countries.

[21] The term 'Black' in South Africa (or in the literature on South Africa) usually refers to all the non-White groups, i.e. Africans, Coloreds, Asians, and others.

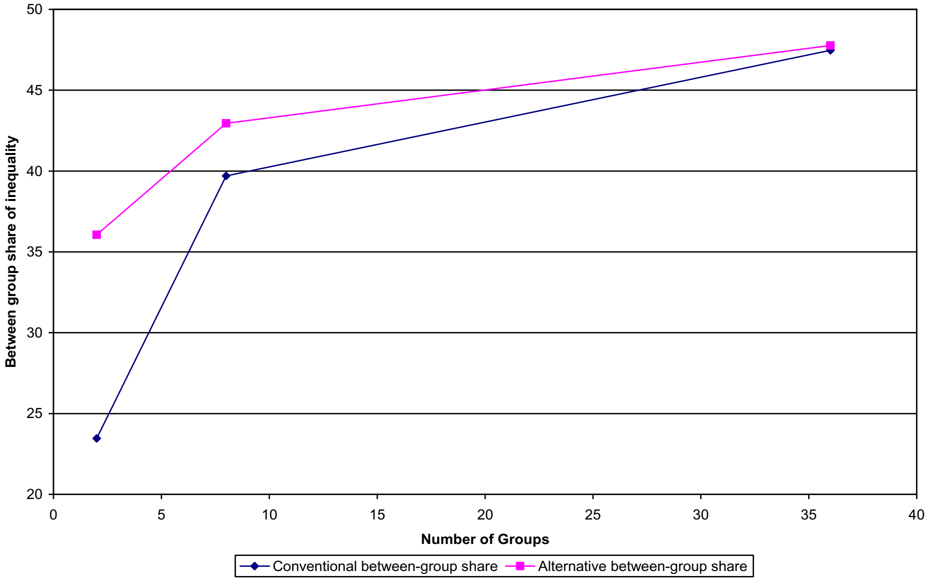

Fig. 2 Finer partitions: urban/rural, then education, then region

Whites and non-Whites? The reason is not only that the per capita monthly mean consumption expenditure of whites (ZAR2210) is much higher than that of non-whites (ZAR407), but also because the range of expenditures for Whites barely overlaps with that of non-Whites (see Fig. 3).[22] On the other hand, while there is a big gap between the mean expenditures of Africans (ZAR357) and non-Africans (ZAR1355), their expenditure distributions overlap quite a bit as the expenditure range of Coloreds and Asians is much closer to Africans than Whites (see Fig. 4).

Of course, the apartheid was about the privileges White people possessed. In South Africa, at least when it comes to economic well-being, it is the "Whites" who are a race apart, not the Africans. The evidence presented above points towards the same conclusion. But, examining inequality between "Whites and non-Whites" as opposed to "Africans and non-Africans" using the conventional inequality decomposition techniques, we would have concluded that these alternative groupings are equally pertinent to our understanding of inequality in South Africa.

Table 3 also shows that moving from two broad sub-groups to the four major racial sub-groups and then further to a finer grouping of race and language, $\widehat{R}_B$ continually declines—eventually to 37%, roughly the same as R_B. This indicates that the differences within 'Blacks' (or the further ethnic differences within the African population for that matter) are much less important in "understanding" inequality in South Africa than those between Whites and non-Whites. If a policy-maker was concerned with racial income inequality in South Africa, it would make much more sense to view the population as White and non-

[22] ZAR stands for South African Rands.

Table 3 Decomposing inequality by ethnic groups in South Africa

	No. of "ethnic" groups	R_B	$\widehat{R}_B$
White/non-White	2	27.1	79.6
African/non-African	2	28.5	49.9
Racial groups	4	33.3	56.4
Racial groups and language	20	36.2	37.4

R_B is the conventional share of between-group inequality in total inequality (measured by mean log deviation in per-capita consumption), while $\widehat{R}_B$ is our proposed measure. There are four major racial groups in South Africa: Africans, Coloreds, Asians, and Whites. 11 major languages are spoken in South Africa. Whites and Coloreds mainly speak Afrikaans and English, Asians overwhelmingly speak English, while all but 1% of Africans speak one of the other languages spoken in South Africa. To avoid groups with very few observations, we combined some of the less widely spoken languages together, yielding 20 groups instead of

White, rather than focusing on the differences between various ethnic groups within the African population.[23]

3.1 Correlating total inequality and between-group inequality

As mentioned in Section 1, Kanbur [10] has cautioned against concluding that simply because (conventionally calculated) between-group contributions to inequality are generally low, this should be taken to imply that between-group differences are of only limited importance to an overall assessment of inequality. As argued most recently in the World Development Report 2006 [16], overall inequality in the developing world tends to be high and to persist over long periods of time in those countries where there exist significant inequalities of opportunity across population sub-groups.[24] Such inequalities may, in turn, act as a brake on economic growth and dampen prospects for rapid poverty reduction. In the spirit of probing this issue further we ask here whether, across a large set of countries, there is any statistical relationship between overall inequality and the percentage contribution that is attributable to between-group differences. Given that our between-group measure, $\widehat{R}_B$, is more readily comparable even where different populations have different group definitions, we employ a cross-country regression framework to study the relationship between overall inequality and between-group differences. Using the country as the unit of analysis, we regress overall inequality separately on the between-group contribution attributable to four population breakdowns: rural–urban location of residence, social group, occupation of household head, and education of household head.[25]

[23] In Elbers et al. [8], we advance the notion that the 'salience' of a particular group to the analysis of inequality may be assessed by examining whether income is a good predictor of membership in that group. In the above example, the success rate in guessing whether one is White or not would be much higher than that in trying to guess whether they are African or not, and still much higher than that of guessing someone's particular ethnic group. $\widehat{R}_B$ seems better suited to fit this notion of salience than R_B. Also see Yizhaki and Lerman [17], pp. 319–320, for a similar discussion.

[24] In what follows, only differences between 'social groups' in these countries can strictly be interpreted as inequality of opportunity in the Roemer sense. The income/consumption differences between other groups, such as rural–urban, education, and occupation are likely due, at least in part, to choices people have made.

[25] Our data come from nationally representative household surveys from each country for a year during the 1990s and are not strictly comparable as inequality is typically measured differently across countries—based sometimes on a consumption measure of welfare and sometimes on an income measure. See World Development Report 2006 [16] for details on the data used in this sub-section. Box 2.5 (p. 38) of the same report provides a more detailed discussion of the issues of data comparability.

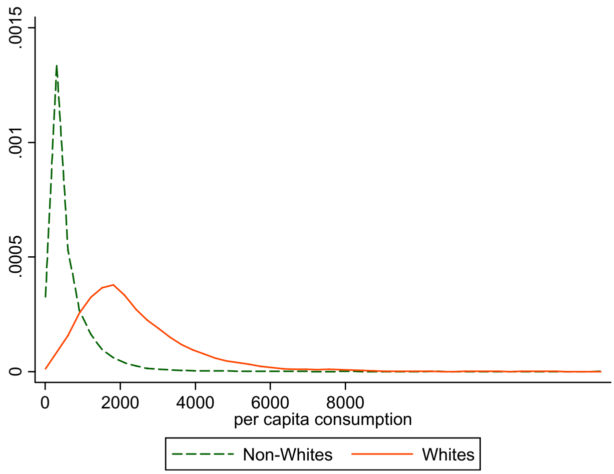

Fig. 3 Probability distribution functions for Whites and others in South Africa

Figure 5 presents our results. We include in our regression a set of regional dummy variables as well as a dummy indicating whether a particular country's inequality is measured on the basis of per-capita consumption or income. Regression results have also been screened for the influence of outliers and influential observations. There is strong evidence of a positive correlation between overall inequality and the between-group contribution, irrespective of the specific group definition. It is important to realize that there is nothing inherent in the mechanics of the decomposition calculation that ensures that there should be a positive relationship between the overall level of inequality and the percentage contribution that can be attributed to between-group differences.[26] We can see that in each case considered here there is a strong and significant correlation between overall inequality and between-group differences.

These correlations are suggestive but, of course, far from conclusive. Nevertheless they are consistent with the above-mentioned arguments advanced in the 2006 World Development Report [16]. To the extent that overall inequality dampens prospects for poverty reduction (given a growth rate), it seems that policy makers have an important reason for concentrating on reducing group differences alongside their possible intrinsic objections to inequality.

4 Concluding remarks

In this paper, we attempt to address two difficulties in interpreting inequality between groups, namely *comparability* and the rather *extreme benchmark* against which between-group inequalities are judged, by proposing an alternative measure. Specifically, we suggest replacing total inequality in the denominator of the conventional ratio with the maximum between-group inequality that could be obtained if the number of sub-groups and their sizes were restricted to be the same as for the numerator. Because our proposed measure is normalized by the number and relative sizes of sub-groups under examination, comparisons are easier across settings where these parameters are very different.

[26] Indeed, if there were concerns about noise in the data, high inequality countries would likely be countries in which there was more noise. Pure noise would result in smaller between-group shares (because of greater overlap across groups). As a result, if anything one might expect a *negative* relationship.

Fig. 4 Probability distribution functions for Africans and others in South Africa

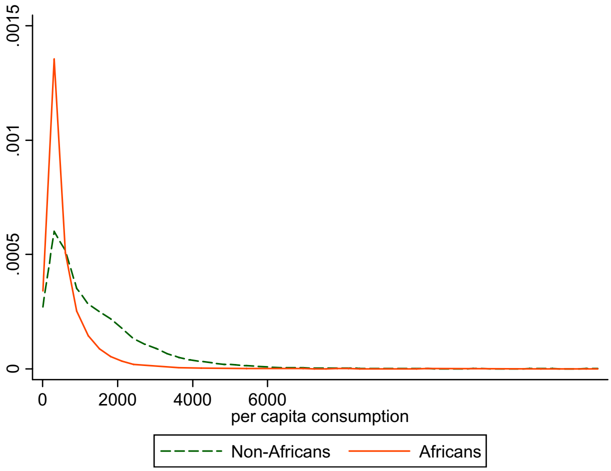

It is important to stress that our measure is not the result of a statistical decomposition exercise for any inequality measure of a certain class. $\widehat{R}_B$ is concerned with evaluating between-group inequality against a proper benchmark and as such places less emphasis on inequality within sub-groups. Our measure is simple to calculate, particularly when we preserve the "pecking order" of the sub-groups under examination.

We suggest that our approach can provide a complementary perspective on the question of whether (and how much) a particular population breakdown is salient to an assessment of

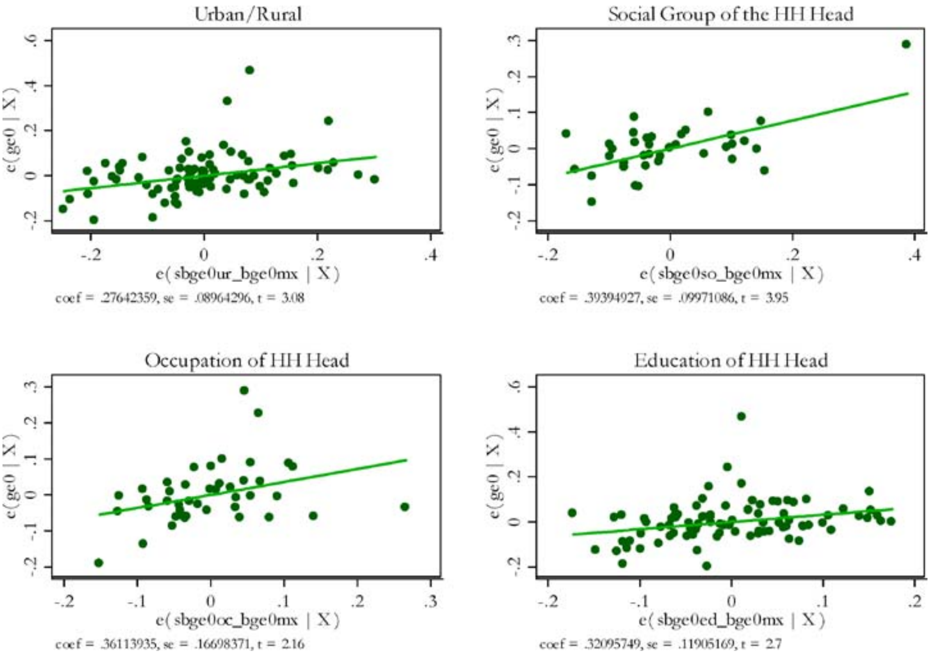

Fig. 5 Regressions of total inequality on shares of between-group inequality of different household characteristics (based on $\widehat{R}_B$)

inequality in a country. Qualitative assessments of the importance of between-group differences can at times be markedly different when based on our alternative approach. For example, if we think of South Africa as comprised of Whites and non-Whites, income inequality between these two sub-groups accounts for 80% of the maximum inequality attainable in South Africa between two such sub-groups (Table 3). The same figure is 50% if we break the population into Africans and non-Africans instead. Interestingly, the conventional decomposition method would have yielded almost exactly the same between-group inequality share (approximately 28%) for these alternative partitions of the population. Viewing South Africa via our alternative approach, a policy-maker concerned with racial inequalities in income would note that the differences across its four major racial groups (56%), or its 20 ethnic groups (37%), pale in comparison to the differences simply between Whites and others.

Acknowledgement We are grateful to Tony Atkinson, Francois Bourguignon, Sam Bowles, Valentino Dardanoni, Jean-Yves Duclos, Francisco Ferreira, Gary Fields, Yujiro Hayami, Ravi Kanbur, Takashi Kurusaki, Peter Lambert, Jenny Lanjouw, Branko Milanovic, Keijiro Otsuka, Adam Przeworski, Martin Ravallion, Tony Shorrocks, and Jacques Silber, participants at a seminar at Foundation for Advanced Studies on International Development in Tokyo, and participants at the First Meeting of the Society for the Study of Economic Inequality in Palma de Mallorca for comments and/or helpful discussions. We are particularly indebted to Marta Menéndez for numerous contributions to this project. We thank the government of Japan, via its Millenium Policy and Human Resource Development Fund Grant, for financial support. The views in this paper are our own and are not to be taken to represent those of the World Bank or any affiliated institution.

References

1. Anand, S.: Inequality and Poverty in Malaysia. Oxford University Press, Oxford (1983)
2. Bourguignon, F.: Decomposable income inequality measures. Econometrica **47**, 901–920 (1979)
3. Cowell, F.A.: On the structure of additive inequality measures. Rev. Econ. Stud. **47**, 521–531 (1980)
4. Cowell, F.A.: Measurement of inequality. In: Atkinson, A.B., Bourguignon, F. (eds.) Handbook of Income Distribution. North Holland, Amsterdam (2000).
5. Cowell, F.A., Jenkins, S.P.: How much inequality can we explain? A methodology and an application to the United States. Econ. J. **105**(429), 421–430 (1995)
6. Cowell, F.A., Kuga, K.: Additivity and the entropy concept: An axiomatic approach to inequality measurement. J. Econ. Theory **25**, 131–143 (1981)
7. Elbers, C., Lanjouw, P., Mistiaen, J.A., Özler, B., Simler, K.: On the unequal inequality of poor communities. World Bank Econ. Rev. **18**(3), 401–421 (2004)
8. Elbers, C., Lanjouw, P., Mistiaen, J.A., Özler, B.: Re-interpreting sub-group inequality decompositions. World Bank Policy Research Working Paper 3687 (2005)
9. Foster J., Sen A.K.: On economic inequality. Expanded edition. Clarendon, Oxford (1997)
10. Kanbur, R.:Income distribution and development. In: Atkinson, A.B., Bourguignon, F. (eds.) Handbook of Income Distribution. North Holland, Amsterdam (2000)
11. Mookherjee, D., Shorrocks, A.: A decompositions analysis of the trend in UK income inequality. Econ. J. **92**(368), 886–902 (1982)
12. Özler, B.: Not separate, not equal: Poverty and inequality in post-apartheid South Africa. Econ. Dev. Cult. Change **55**(3) 487–529 (2007)
13. Shorrocks, A.F.: The class of additively decomposable inequality measures. Econometrica **48**, 613–625 (1980)
14. Shorrocks, A.F.: Inequality decomposition by population subgroups. Econometrica **52**, 1369–1385 (1984)
15. Shorrocks, A.F., Wan, G.: Spatial decomposition of inequality. J. Econ. Geogr. **5**, 59–81 (2004)
16. World Bank: World development report. Equity and development. Oxford University Press, Washington D.C. (2006)
17. Yitzhaki, S., Lerman, R.I.: Income stratification and income inequality. Rev. Income Wealth **37**(3), 313–329 (1991)