From social choice to inequality-decomposition: In the spirit of Arrow and Atkinson by way of Sen and Shorrocks*

Authors

Ben Fine. Emeritus Professor of Economics at the School of Oriental and African Studies, University of London, UK, and Visiting Professor, Wits School of Governance, University of Witwatersrand, South Africa. Email: bf@soas.ac.uk. ORCiD: https://orcid.org/0000-0002-4301-0730.
Pedro Mendes Loureiro. Lecturer in Latin American Studies at the Centre of Latin American Studies and Department of Politics and International Studies (CLAS-POLIS), University of Cambridge, and Fellow of Fitzwilliam College, University of Cambridge. Email: PML47@cam.ac.uk. ORCiD: https://orcid.org/0000-0001-7517-7351. Corresponding author.

Abstract

This article forges connections between social choice theory and inequality-measurement to deliver a series of advances to the latter. It discusses the ethical and formal aspects connecting the fields, linking the roles played by the intensity of preferences and interpersonal comparability (in social choice) and, respectively, inequality-aversion and interpersonal comparability (in inequality-measurement). This extends naturally to relaxing the assumption of symmetry in group-wise inequality-decompositions, allowing a measure capable of unambiguous decomposition that, overcoming well-known limitations of existing measures, still incorporates different degrees of inequality-aversion through use of its demonstrated duality with inter-group comparability. The framework is then applied to inequality in Brazil and the changes it underwent between 2003 and 2015, focusing on labour-market incomes of white men and non-white women. We show how sensitive are measures of (changes in) inequality to assumptions regarding the interpersonal comparability of the groups' income and inequality-aversion, pinpointing the precise trade-offs between the two. The results indicate that, for a reasonable range of parameters, overall

*Arrow died on February 21, 2017 (aged 95), and Atkinson on 1 January 2017 (aged 72).

inequality changed by between -28.6% and -12.2% over the period. This highlights the analysis of the trade-off between interpersonal comparability and inequality-aversion warrants being addressed explicitly in empirical measurements of, or judgements over, levels and variations of inequality.

Keywords

Inequality Measurement; Social Choice; Inequality Decomposition; Brazil; Ethnoracial Inequality; Gender Inequality.

1 Introduction and Preliminaries

There has been a longstanding if shifting relationship between social choice theory and measurement of inequality. Although social choice thrust to prominence with Arrow's (1951) Impossibility Theorem, it is based upon a generalisation of the voting paradox, traditionally associated with the Marquis de Condorcet at the end of the eighteenth century. Measuring inequality has been more prominent for longer, not least with the ubiquitous Gini, invented by Corrado Gini in 1912, drawing upon Max Lorenz's curve comparing cumulative distribution from the bottom up to perfect equality.[1]

The two approaches share some obvious common elements, both formal and informal. Each seeks at least to rank, possibly to quantify, a set of alternatives (for society or for the more specific case of distributions of income). In doing so, each inevitably makes at least implicit, and often explicit, ethical judgements, as rankings tend to carry the implication of being for better or for worse. And, especially over the past sixty years, economists and others have been applying mathematical methods to explore the relationship between desirable properties of their social choices or inequality measures and their consequences. The strongest intersection between the two approaches has been forged by Amartya Sen, a major, if rare, contributor to both fields, not least as part of his intellectual trajectory began with social choice, moved to inequality and, from there through entitlements and capabilities, to freedom – see Sen (1970a, 1973 and 1999) and Fine (2004).[2]

Nonetheless, there are also significant differences between the fields that have tended to send them along parallel routes from time to time. Social choice is more geared towards the ethical (or desirable property) side of things, not least seeking to manoeuvre around the deadweight of Arrow's founding Impossibility Theorem. It is hard to think of empiri-

[1]Arrow is well-known for wide-ranging contributions to mainstream economics, especially general equilibrium theory, receiving the Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel in 1972, with John Hicks, in its fourth year; Gini had strong fascist leanings; Condorcet, although a supporter, died in mysterious circumstances in the wake of the French revolution; and Lorenz was a statistician in the US Government.

[2]See also Fine (2018) for a (re-)assessment in light of republication of Sen (1970a) with as much new material as old, Sen (2017). Perhaps Sen is an exception in straddling the two fields at the highest levels, with Arrow contributing to social choice theory but not inequality measurement, and vice-versa for Atkinson (and Shorrocks, see below).

cal or practical applications of social choice theory although such social choices take place in a formal sense all the time, not least in elections and voting more generally. Measures of inequality (essentially variance of income) have been much more pragmatic, more concerned with statistical than ethical properties – just tell me the Gini and whether it has changed or not, with growth or otherwise for example. Nonetheless, the exploration of the ethical principles underpinning social choice theory, not least with their mathematical formulation and investigation, has significant implications for the statistical measurement of inequality that have tended to be overlooked.

The purpose of this paper is to offer some remedy by knitting the two fields together in such a way that important advances can be made in measuring inequality in various ways (as well as pointing to reasons why these advances have not been made previously because of the failings reflected in integrating the two fields). To this end, Section 2 offers an overview of some key, relevant aspects of social choice theory and its motivation before doing the same for inequality measurement in Section 3, drawing in part upon, and briefly presenting new formal results derived from, Fine and Loureiro (2020). In particular, Section 3 reviews the relationship between welfare-motivated measures of inequality, especially associated with Atkinson (1970), and direct measures of income inequality, associated with Shorrocks (1980). More precisely, as is already known if far from prominently highlighted, Atkinson's requirement that *social welfare* be additively separable (it is the sum of individual welfares) and homothetic (increasing if unevenly for the same proportionate distributions) is shown to be slightly more stringent than Shorrock's requirement that the measure of income *inequality* be decomposable (capable of disaggregation of overall income into a sum of within and between sub-group inequalities) and homogeneous of degree 0. However, the source of that difference has been previously identified technically (in the respectively derived differential equations to be solved) and in the addition of an overall income term in the measure of welfare for Atkinson (Fine and Loureiro 2020).

Such a resolution between the approaches of Atkinson (to income inequality derived from aggregate welfare) and Shorrocks (to measure income inequality directly) would appear to be of relatively minor significance in and of itself. It does, however, set up a duality between the two approaches of considerable interpretative and even practical importance. This is especially so, and revealed as such through the novel results presented in Section 4, once the axioms underpinning the respective measures are relaxed, with symmetry in particular in mind (for which each individual is treated as the same or each is substitutable for one another without affecting the choice/measure). Dropping symmetry also has the virtue of reflecting particular motives for measuring inequality, and its decomposition, in the first place, especially when we are less concerned with individuals as such and more with their shared properties or as groups that experience some degree of (dis)advantage in common. Specifically, do we wish to treat men the same as women in measuring (changes in) inequality or similarly for different racial or ethnic groups. Precisely by dropping symmetry for Atkinson's approach, we find that the measure for inequality allows for interpersonal comparisons (one person's income/welfare counts more or less than another's) not only in a natural way but also one that is equivalent to a relatively straightforward transformation of the income distribution itself – we multiply each individual income by a weight. Through the previously observed duality, these transformed incomes can then be deployed within the Shorrocks' measure (and for decomposition of inequality) as if they were actual incomes.

As such, this still does not seem to gain any advantages; you can measure in one system or the other. But, unsurprisingly, in the absence of symmetry the Atkinson measure means that weighting individual welfares (or their incomes) more or less is equivalent to being more or less inequality-averse, as you weight the poor or rich more or less. This too carries over into the Shorrocks measure. Here, though, there is a crucial added advantage to be found, around addressing the issue of decomposability – this involves disaggregating overall inequality into that contributed by the partition into subgroups, with the inequality assigned to the appropriately weighted sum of inequality within each group and the inequalities between groups. Conditions for attaining such decomposition is Shorrock's explicit aim but, in his closing discussion, he observes there is a problem with his measures since weights for within sub-group inequality do not in general add up to one, and their sum is not independent of the between sub-group measures. In addition, there is some ambiguity in the way in which inequality between sub-groups is specified – do we simply compare inequality as if each different sub-group had no inequality (within group is zero) and the only source of inequality is difference in the sub-group means; or do we compare for overall distribution as if each sub-group has the same mean but with their intra-distributions otherwise remaining the same.[3]

Shorrocks then observes that these conundrums are eliminated, and weights sum to one and the different ways of comparing decompositions become the same, in case just one of the available functional forms for measuring inequality is adopted, namely, as if total welfare is $\sum \log x_i$ for individual incomes, x_i, in Atkinson terms. For Shorrocks, this is where an inequality-measuring parameter, c, takes on the value 0,[4] for which he concludes, p. 625:[5]

> For this reason, [this] is the most satisfactory of the decomposable measures, allowing total inequality to be unambiguously split into the contribution due

[3]Difficult to express but, as Shorrocks (1980, p. 624) puts it himself (see Deutsch and Silber 1999 for a longer discussion of how these principles apply to different decompositions of inequality):

> Interpretation (i) suggests comparison of total inequality with the amount which would arise if inequality was zero within each age group, but the difference in mean income between age groups remained the same. For the additively decomposable indices this would eliminate the total within group term and leave only the between group contribution ... Interpretation (ii) suggests a comparison of total inequality with the inequality value which would result if the mean incomes of the age groups were made identical, but inequality within each age group remained unchanged.

[4]Only for Shorrocks' $c = 0$ is the contribution of each group to within-groups inequality a *population*-weighted average of their intra-group inequality indices, as opposed to $c \neq 0$ when relative incomes enter into the weights of the summation. As such, $c = 0$ is the only level of inequality-aversion in which within-groups inequality is truly independent of the between-groups dimension.

[5]Of course, it is possible to continue to use other measures that do not have the properties of the log form but at the expense of reintroducing the highlighted ambiguity in decompositions. See below for full discussion.

> to differences between sub-groups ... and the contribution due to inequality within each sub-group.

However, this is not entirely satisfactory because of two closely related issues, each of which can be resolved as a result of the duality between the approaches of Atkinson and Shorrocks that have been forged. First, by virtue of decomposing inequality by sub-groups, whether specified formally through an anonymous index or concretely by some social characteristic, asymmetries are being explicitly presumed. Why otherwise would we want to decompose whether by age, gender or some other aspect? As already mentioned previously, we are treating sub-groups as different from one another. As a result, we lose the rationale for retaining the axiom of symmetry given individuals, and the groups to which they belong, are not the same for reasons other than income alone (although lower or higher income is often reflected in such group differences and is a rationale for decomposing into groups at the outset). In short, there are strong motivational reasons for rejecting the axiom of symmetry in undertaking exercises in decomposability.

Second, in order to address the statistical conundrums identified by Shorrocks (that in decomposing overall inequality into a weighted average of intra-group inequalities), a single measure of inequality-aversion must be adopted (the log form). However, in the absence of symmetry, we have alternative options for, indirectly, incorporating inequality-aversion (or even inequality-preferring). We can assign different income/welfare weights to different individuals and/or sub-groups and these have the effect of varying the inequality-aversion that would otherwise be absent. This is the beauty of taking the results from Atkinson to their dual in Shorrocks! By doing away with symmetry, we can use the ideal functional form for decomposing inequality but allow for different degrees of inequality-aversion indirectly (since it is directly fixed by the selected $c = 0$ or log form) by introducing interpersonal weights on incomes through relaxation of the symmetry condition (previously adopted by both Atkinson and Shorrocks).

In a nutshell, our paper neatly brings together: first, the exact nature of the established duality between Atkinson and Shorrocks in measuring inequality; second, the consequences of weakening the assumption of symmetry (e.g. men count the same as women) in both a formal sense (what difference does it make to measures) and an ethical sense (the derived duality between intra- and inter- personal judgements); and, third, the use of these results to establish a way to perfect decomposability in measuring inequality without being hamstrung by the necessity of relying upon a single parameter value for inequality aversion. On this basis, by the end of Section 4, by relaxing the assumption of symmetry, it is possible to lay out formulae for measuring inequality that include both inequality-aversion and interpersonal comparability and, in particular, to deploy the trade-offs between them to allow for a decomposable measure of inequality that is free of ambiguity. On the basis of these formulae for measuring and disaggregating inequality, we undertake in Section 5 a select empirical exercise on Brazilian inequality to illustrate the power of the method outlined. We show how sensitive are measures of (changes in) inequality to assumptions over interpersonal comparability of groups, b, and inequality-aversion, c, but also construct iso-inequality and iso-changes-in-inequality curves that show exactly how changing the value of the two parameters, b and c, are related to one another. Significantly, if unknowingly and inconsistently, some may find an inequality-aversion parameter acceptable but not its equivalent comparability parameter, an issue taken up in the final remarks.

2 From Social Choice ...

Social choice theory effectively began with Arrow's (1951) classic contribution. It had no obvious prominent predecessors and reflects Arrow's own preoccupation with soundly constructing a social welfare function. As such, it was very much a product of its time at a number of different levels.[6] The first, and most general, was to create the foundations for making public choices in a context, as the post-war period began, of increasing state interventionism.

Second, and surely not accidental given Arrow's even more prestigious contributions to general equilibrium theory, the goal became one of founding social choice as if the state were an individual with corresponding preferences of its own over available alternatives. Under what circumstances would political choices on the basis of individuals work out well, just as the same might, or might not, be realised through competitive market forces for the economy. In this respect, like Solow for the aggregate production function, Arrow was initiating an approach in which a way was sought to extrapolate from an individualistically-based ethos to a social one. However, unlike Solow, for whom the difficulties involved needed to be teased out by others in the Cambridge Capital Controversies,[7] Arrow's contribution to social choice brought out inherent tensions from the outset, with his impossibility theorem. It would appear that you could neither aggregate preferences nor capital, albeit for very different (technical and conceptual) reasons![8] In short, no social choice in general means no implications for measuring inequality in particular.

Third, such reductionism of social to (unexplained and unexplored) individual preferences over available alternatives allowed for another gathering ethos of the time to be adopted, namely mathematical formalism. Arrow's, and subsequent contributions, have been heavily oriented around the consequences of axioms on choice mechanisms. Indeed, this is the elegance of Arrow's contribution, equally setting the pattern of what was to come. With four apparently innocuous axioms – no dictator, completeness (with derived choice required from any set of individual preferences), all feasible sets of preferences to be

[6] For a more rounded discussion of 'welfarism' and its position in the history of economic thought, see Backhouse et al (2020).

[7] See Fine (2016). Consider how limited has been critical attention to social choice as opposed to the Cambridge capital controversy, and how different things might have been if General Equilibrium had been posed as an impossibility rather than an existence theorem in light of the restrictions necessary for it to hold.

[8] Significantly, although neither has ever subscribed to the conceptually crude forms it has taken with Becker, both Arrow and Solow can be seen as contributing to economics imperialism, seeking to treat the society and economy as if a single individual (at the expense of social relations, structures, etc), with game theory also emerging out of the state as individual engaging in Cold War nuclear deterrence. See Fine and Milonakis (2009).

accommodated (unrestricted domain), and each social choice between pairs of alternatives being consistent with binary choice over them (independence of irrelevant alternatives) – Arrow demonstrated that no social choice could exist.

Having got this result, it was much more natural to seek out a weakening of the axioms than to add others, not least that might allow for addressing the specifics of inequality – after all, impossibility was already there and so further conditions could not be accommodated. The position pursued here is that the most constraining, even unreasonable axiom, is that of irrelevant alternatives, without which potential for more axioms and reasonable choices are opened up. This can be seen in two ways. First, suppose I prefer an alternative $\boldsymbol{a}$ more than a million other things and prefer these million over $\boldsymbol{b}$ as well. Surely this is relevant in social choice between $\boldsymbol{a}$ and $\boldsymbol{b}$ compared to the situation in which my preference between the two alternatives is barely marginal, and no other alternative comes between $\boldsymbol{a}$ and $\boldsymbol{b}$? Such a situation says something about the intensity of preference between alternatives. To anticipate, it has some resonances with how much we might appreciate more income rather than just seeing more as more never mind by how much and how much valued (by the individual or those assessing the worth for the individual in making inequality measures).

Second, Arrow's result, which is a grand generalisation of the voting paradox, depends upon some unpalatable consequences as a result of the irrelevance axiom. With individuals 1, 2 and 3, and alternatives $\boldsymbol{x}$, $\boldsymbol{y}$ and $\boldsymbol{z}$, the paradox can be represented as follows, with intransitive majority rule of $\boldsymbol{x}$ over $\boldsymbol{y}$ over $\boldsymbol{z}$ over $\boldsymbol{x}$, by two to one.

1	2	3
$\boldsymbol{x}$	$\boldsymbol{y}$	$\boldsymbol{z}$
$\boldsymbol{y}$	$\boldsymbol{z}$	$\boldsymbol{x}$
$\boldsymbol{z}$	$\boldsymbol{x}$	$\boldsymbol{y}$

But simple inspection of the preference structure across individuals and alternatives shows them to be completely symmetrically placed in relation to one another, or isomorphic (if not by binary choice, i.e. setting aside $\boldsymbol{z}$ when considering $\boldsymbol{x}$ and $\boldsymbol{y}$, for example). So, the only sensible way in which to come to a social choice is to treat each alternative (and individual) as the same. So none should win and social choice should be total indifference (in the formal sense - there is nothing by which to choose between the alternatives rather than not caring or being unable to make a choice).

This is the starting point taken by Fine and Fine (1974a and 1974b) and Fine (1974). It has the advantage of not only dissolving the impossibility but also of creating the space for other, and most reasonable, axioms to be considered, with considerable knock-on effects for the measurement of inequality.[9]

As a result, a more constructive approach is made in the following way. First is to assume that there is total social indifference for isomorphic sets of preferences (i.e. if swapping names of individuals and/or alternatives makes no difference to the preference structure, there should be social indifference as for voting paradoxes). Second is to adopt

[9] For a recent but different approach to producing partial rankings, discussed in the context of inequality measurement, see Cato (2020).

the axiom of monotonicity. If, within a configuration of preferences, an alternative $\boldsymbol{x}$ is preferred to (or equal to) $\boldsymbol{y}$ and, in another configuration, the only difference is that $\boldsymbol{x}$ is moved up or $\boldsymbol{y}$ is moved down, then $\boldsymbol{x}$ must remain preferred (or become preferred to) $\boldsymbol{y}$. Third is to allow for composition (this is close to decomposability for measuring inequality, see below). If $\boldsymbol{x}$ is preferred to $\boldsymbol{y}$ in two sets of individuals with preferences over a given set of alternatives, then so it should remain for the combined group of individuals.

With these axioms alone, it is possible to restrict the set of choices that can be made extremely narrowly. The way to do so is to take any configuration of preferences and decompose them into a number of sets of isomorphic preferences that are adjusted for monotonicity. In this way, we move from indifference to positive outcomes, and social choice will depend on a few basic set of preferences alone, in which whether alternatives come first, second or third, etc, for individuals, can be traded off against one another. More precisely, social choice is reduced to what are termed finite ranking rules, given by a vector $(c_1, c_2 \ldots c_n)$ for n alternatives with $c_1 > c_2 > \ldots > c_n$ in which c_i is awarded to each alternative for an i-th ranking by an individual and the alternatives are ranked according to their aggregate scores across all individuals.][10]

More axioms could be added to restrict the values of the c_i. For example, for a weaker version of irrelevant alternatives – that outcomes between alternatives $\boldsymbol{x}$ and $\boldsymbol{y}$ should not depend on alternatives that are ranked by all individuals either above or below $\boldsymbol{x}$ and $\boldsymbol{y}$ (but never between $\boldsymbol{x}$ and $\boldsymbol{y}$), then the c_i increase in geometric progression.[11] More generally, the c_i represent some sort of intensity of individual preference, quantifying how much weight is given to an alternative should it move from a lower to a higher position. This has some resonances with measurement of income inequality as more or less weight is given to higher incomes. More generally, restrictions might be placed on the range of values to be taken by the c_i and, in this case, some partial ordering of the alternatives will result that will become definite at least by the time a single vector for the c_i is fixed. Otherwise, for example, the alternative with the most first positions can never be beaten however badly it might perform since a large enough c_1 relative to the other c_i will always allow it to win, Fine (1996).

So far, for reasons that should be apparent, social choice has been discussed in terms of symmetry of individuals (and alternative choices). Each is treated the same as any other. But there may be a case for violating such symmetry in a way that parallels the restrictions just discussed for the intensity of preference, as suggested by Sen (1970b). More specifically, one way to treat asymmetry is to do so through interpersonal comparability by weighting each individual's given utility by more or less than others, by a parameter, b_i, say, for individual i. This, for example, might get us out of rigid adherence to the Pareto principle for which just one individual with marginally higher welfare than anyone else would block a move in which everyone else got loads more but this individual marginally less.[12] As long as we have a reasonable, but not full, range of the parameters

[10] For a more general discussion of such positional rules, see Bossert and Suzumura (2020).

[11] If, in addition, choice is reversed when preferences are reversed, then the Borda rule results (equal gap in rankings in moving up preferences). Note this is all about the intensity of preference in moving up or down the individual rankings, corresponding to inequality aversion in welfare measurement.

[12] And, of course, from Lionel Robbins onwards, Pareto has been death to the ethical measurement of

b_i, it becomes possible to decide against extreme application of Pareto efficiency as sole criterion for positive choice.

Thus, in moving to social choice, as argued by Fine (1996), there are two avenues to take, either sequentially or in parallel: one is to restrict intensity of preference, and the other is to restrict extent of interpersonal comparison, the c_i and b_i, respectively. It is also shown that these are to some degree equivalent to one another, with greater restrictions on the intensity of preference (i.e. not allowing enormous increases in welfare to individuals who are already well off) equivalent to favouring the welfare of those who are not so well off. This is hardly surprising but it is important to acknowledge that they have different origins, or ethical foundations, even if leading to similar outcomes of judgement. For one is about intra-personal welfare (how much better am I as I move up my preferences) as opposed to the other being about how much does one person's movement up count against somebody else's. Exactly the same considerations apply to the measurement of income inequality since we are dealing with how much extra income is worth to an individual in and of itself (intensity of preference) and by comparison with income to another (interpersonal comparison). This is why social choice freed from the chains of Arrow's impossibility theorem (and the logically and ethically unacceptable axiom of irrelevant alternatives) is of such potential importantance to the measurement of inequality.

3 ... To Measuring Inequality

As a starting point, it can be observed that there are two distinct but closely related approaches to the theory of constructing measures of income inequality from first principles. One is indirect, proposing more general principles for social welfare, based on income, from which implications for inequality can be deduced. Unsurprisingly, this is the approach that emerged most strongly in the wake of social choice theory, especially after Sen (1970a and 1973) and, with its dependence on axioms, has strong affinities with social choice. The other approach is to jump the first stage and, as a special case of it, engage with inequality directly by drawing out consequences from first principles targeted on income distribution. It is certainly not ethically neutral in the sense of merely concerning itself with a measure of variance but it is more inclined in that direction. Both of these approaches, however, contrast with the more pragmatic approach of adopting and using a measure with greater or lesser regard for its properties that may or may not be uncovered and critically scrutinised for their ethical content.

A classic example of the first approach is provided by Atkinson (1970). He assumes that social welfare is additively separable to draw out the implications for inequality. This means that welfare W is simply given by the sum of individual welfares. This might be justified in the context of concern with income alone, isolated individuals, and no reason available other than to sum individual welfares to assess welfare. In addition, Atkinson assumes that the measure for W is homothetic, meaning that the rate at which it changes between the individuals is invariant to the ratios of indidvidual incomes. Homotheticity is

inequality given prohibition of interpersonal comparisons.

a weaker condition than homogeneity with the scale at which W changes with the vector of incomes not fixed by some degree of homogeneity.

On this basis, Atkinson essentially shows that, if a corresponding measure of inequality is standardised for overall income, for it to lie between 0 (perfect equality) and 1 (perfect inequality), and with n individuals, the only measures possible that satisfy the conditions are:

$$I = 1 - \left[\frac{1}{n}\sum\left(\frac{x_i}{\mu}\right)^{1-\epsilon}\right]^{\frac{1}{1-\epsilon}}, \epsilon \geq 0, \epsilon \neq 1$$
$$I = 1 - \exp\left[\frac{1}{n}\sum\log\left(\frac{x_i}{\mu}\right)\right], \epsilon = 1 \tag{1}$$

where x_i is the i-th individual's income, μ is the mean income and ϵ is the inequality-aversion parameter (with higher ϵ indicating higher aversion).[13]

What this all reveals, however, is just how powerful are the principles of additive separability and homotheticity in restricting the functional forms that can be taken in measuring overall welfare, with corresponding implications for inequality. The same follows from the other approach, addressing the theory of measuring inequality directly, as illustrated by Shorrock's (1980) classic piece in which the target is one of decomposing measures of inequality into within sub-group and between sub-group contributions. This is an attractive target for allowing assessments of how much each sub-group contributes to inequality and how much does inequality between sub-groups (the reader is referred to the original contribution for details).

Let $\boldsymbol{x}$ be a vector of incomes for a population of size n. Let the population be reordered at will (a way of defining symmetry), and partition the population into sub-groups of vectors of incomes $\boldsymbol{x}_1, \boldsymbol{x}_2, \ldots, \boldsymbol{x}_m$, in m sub-groups, each of size n_g, with $n = \sum_g n_g$. Decomposability is essentially defined by Shorrocks to mean that our measure of inequality, I, a function of the incomes and population size, takes the following form for any partition (division of the population into mutually exclusive sub-groups) $\boldsymbol{x}$, where μ_i are partition sub-group means and w_i weights:

$$I(\boldsymbol{x}, n_x) = \sum_g \left(w_g I(\boldsymbol{x}_g, n_g)\right) + I(\mu_1, \ldots, \mu_1, \mu_2, \ldots, \mu_2, \ldots, \mu_m, \ldots, \mu_m, n_x) \tag{2}$$

Decomposabililty simply means we can disaggregate overall inequality across any set of sub-groups of the incomes into two parts. One is to take a weighted, w_g, combination of the inequalities from the inequality measures of the sub-groups. The other is the inequality measure as if each sub-group had equally distributed mean of its income. Taking I to be lower bound by zero for perfect equality and assuming I is symmetric in incomes, and

[13] ϵ is used to conform with Atkinson's (1970) notation and to distinguish it from the inequality-aversion parameter c, in Shorrock's framework, which as discussed below has a different meaning even if measures in one are ordinally equivalent to those in the other.

that I is homogeneous of degree 0 in incomes, without any other assumptions, Shorrocks shows that I necessarily takes the functional forms (when not decomposed but satisfying decomposability), as follows:[14]

$$\begin{aligned} I(\boldsymbol{x}, n) &= \frac{1}{n\mathrm{c}(\mathrm{c}-1)} \sum_i \left[\left(\frac{x_i}{\mu} \right)^{\mathrm{c}} - 1 \right] && \text{for c} \neq 0, 1 \\ I(\boldsymbol{x}, n) &= \frac{1}{n} \sum_i \log \frac{\mu}{x_i} && \text{for c} = 0 \\ I(\boldsymbol{x}, n) &= \frac{1}{n} \sum_i \frac{x_i}{\mu} \log \frac{x_i}{\mu} && \text{for c} = 1 \end{aligned} \tag{3}$$

where c is the inequality-aversion parameter.[15] It should also be noted that Shorrocks not only derives the functional forms for an overall inequality measure (for the whole distribution) but also uncovers the heavily constrained forms that must be taken by the weights when decomposing the measure. These are shown to be:

$$\begin{aligned} w_g &= \left(\frac{n_g}{n} \right) \left(\frac{\mu_g}{\mu} \right)^{\mathrm{c}} && \text{for c} \neq 0, 1 \\ w_g &= \left(\frac{n_g}{n} \right) = p_g && \text{for c} = 0 \\ w_g &= \left(\frac{n_g}{n} \right) \left(\frac{\mu_g}{\mu} \right) = s_g && \text{for c} = 1 \end{aligned} \tag{4}$$

Decomposition as such is not on Atkinson's agenda but, given symmetry and homogeneity of Shorrocks' inequality function, I, additive separability (à la Atkinson) is a stronger condition than decomposability (à la Shorrocks) since it allows for fewer functional form amongst those that are already remarkably few, although the differences are tantalisingly marginal. But what exactly is it that makes the two conditions different?

As shown by Fine and Loureiro (2020), the difference boils down to subtracting an extra term from the Atkinsonian measure of welfare, namely a function of mean income (indeed, simply the same function that is applied to the individual incomes that are aggregated in Atkinson's original formulation). In effect, Atkinson and Shorrocks coincide if and only if the Atkinson measure of welfare, with homotheticity, is modified to allow for the incorporation of the deduction of a corresponding function of mean income.[16]

[14] There are also some minor technical assumptions which we do not cover.

[15] Differently from ϵ, c is not bound from below by zero. It is more sensitive to bottom-end inequalities (i.e. the presence of low relative incomes) for lower values of c and to top-end inequalities (i.e. the presence of high relative incomes) for high values of c.

[16] All the technical proofs are contained in Fine and Loureiro (2020). Note, though, that the process of obtaining the results reveals more fully the nature of the issues involved through the derivations and proofs involved.

4 Engaging Asymmetry

Now suppose that the Atkinson measure does not satisfy symmetry, treating all individuals the same but, instead, treats them differently whilst retaining additive separability and homotheticity. A natural way to do this, pursued by Fine (1985), is to adopt different functional forms for each individual so that $W = \sum_i f_i(x_i)$.[17] What Fine shows is that the same individual functional forms for symmetry, with parameter, epsilon, remain exactly the same without symmetry except that they can at most differ by an individual scalar parameter. So $W = \sum_i b_i f_i(x_i)$.

Exactly the same applies in case the modified Atkinson welfare measure (i.e. incorporating a deduction of a function of mean income) is adopted with asymmetry where $W = \sum_i f_i(b_i x_i) - g\left(\sum_i b_i x_i\right)$ (see Fine and Loureiro 2020) . g will be a linear transformation of the common f. Further, when converting to a measure of inequality with zero for complete equality, the results for the relationship between Atkinson and Shorrocks will hold on the transformed incomes. For Shorrocks, the transformations proceed as follows, where (upon dropping the subscript i) the transformed coefficients translate (additively separable) welfare into additively separable income inequality for which we subtract $f(\mu)$ from each of the terms for individual income x_i before summing to give the inequality index.

Thus, for Atkinson to become Shorrocks, across the latters' three cases of the functional forms, the measurement of within sub-group inequality remains unaltered as the b_i only serve to make comparisons between different groups. The weights for the contribution of these within sub-groups inequalities are given, for sub-groups $(\boldsymbol{x_1}, \boldsymbol{x_2}, \ldots, \boldsymbol{x_g}, \ldots, \boldsymbol{x_m})$, by:

$$
\begin{aligned}
w_g &= \left(\frac{n_g}{n}\right)\left(\frac{\mathrm{b}_g \mu_g}{\mu^*}\right)^{\mathrm{c}} = p_g \eta_g^{\mathrm{c}} && \text{for c} \neq 0, 1 \\
w_g &= \left(\frac{n_g}{n}\right) = p_g && \text{for c} = 0 \\
w_g &= \left(\frac{n_g}{n}\right)\left(\frac{\mathrm{b}_g \mu_g}{\mu^*}\right) = p_g \eta_g = s_g && \text{for c} = 1
\end{aligned} \tag{5}
$$

where $\mu^* = \sum_g p_g \mathrm{b}_g \mu_g$ is the revised mean for transformed incomes, $p_g = \frac{n_g}{n}$ is the population-share of each group, $\eta_g = \dfrac{\mathrm{b}_g \mu_g}{\sum_h p_h \mathrm{b}_h \mu_h} = \dfrac{\mathrm{b}_g \mu_g}{\mu^*}$ is the (revised) relative income of the g-th group, and $s_g = p_g \eta_g$ is the g-th group's share of (revised) income, with weights only summing to 1 for $\mathrm{c} = 0, 1$ (respectively, weighted by population- and income-shares). The between-group inequality in the decomposition is in turn measured by:

[17] Another, more general way to relax symmetry is by rescaling individual income before forming separably additive welfare, with different individual welfare functions as before. It can be shown this makes no essential difference since the rescaling of income can be incorporated into the different individual functional forms.

$$
\begin{aligned}
I^B &= \left[\frac{1}{n\mathrm{c}(\mathrm{c}-1)}\right]\sum_g \left[n_g\left[\left(\frac{\mathrm{b}_g\mu_g}{\mu^*}\right)^{\mathrm{c}} - 1\right]\right] && \text{for } \mathrm{c} \neq 0,1 \\
I^B &= \left[\frac{1}{n}\right]\sum_g \left[n_g \log\left(\frac{\mu^*}{\mathrm{b}_g\mu_g}\right)\right] && \text{for } \mathrm{c} = 0 \\
I^B &= \left[\frac{1}{n}\right]\sum_g \left[n_g \left(\frac{\mathrm{b}_g\mu_g}{\mu^*}\right)\log\left(\frac{\mathrm{b}_g\mu_g}{\mu^*}\right)\right] && \text{for } \mathrm{c} = 1 \qquad (6)
\end{aligned}
$$

As indicated in the introduction, the ideal for measuring decomposability for lack of ambiguity of interpretation is given by $\mathrm{c} = 0$, with variability in aversion to income inequality available through judicious choice of the interpersonal comparability parameters, b_g. In other words, before we measure inequality, we first adjust within sub-group incomes by a common factor, b_g, and the bigger this factor the more we consider the group to be favoured (in the sense that it derives more welfare from a given unit of income, as detailed below, and so will be treated as if having a higher income than it does). This is not arbitrary but an inevitable consequence of relaxing symmetry which is itself a special assumption in which weights on all groups, or even individuals, are arbitrarily taken to be the same.[18]

The overall measure of inequality and its decomposition into sub-groups can thus be expressed, with some re-arrangements, as:[19]

$$
\begin{aligned}
I_{\mathrm{c}} &= I_{\mathrm{c}}^W + I_{\mathrm{c}}^B = \sum_g w_{\mathrm{c},g} I_{\mathrm{c},g} + I_{\mathrm{c}}^B \\
I_{\mathrm{c}\neq 0,1} &= \frac{1}{n}\frac{1}{\mathrm{c}(\mathrm{c}-1)}\sum_i \left(\eta_i^{\mathrm{c}} - 1\right) = \sum_g \left(p_g \eta_g^{\mathrm{c}}\right) I_{\mathrm{c},g} + \frac{1}{\mathrm{c}(\mathrm{c}-1)}\sum_g \left[p_g\left(\eta_g^{\mathrm{c}} - 1\right)\right] \\
I_0 &= -\frac{1}{n}\sum_i \left(\log \eta_i\right) = \sum_g \left(p_g\right) I_{0,g} - \sum_g \left(p_g \log \eta_g\right) \\
I_1 &= \frac{1}{n}\sum_i \left(\eta_i \log \eta_i\right) = \sum_g \left(s_g\right) I_{1,g} + \sum_g \left(s_g \log \eta_g\right) \qquad (7)
\end{aligned}
$$

where I_{c} is the inequality index for parameter c, I_{c}^W is the within-groups component, I_{c}^B is the between-groups component, $w_{\mathrm{c},g}$ is the weight of the g-th group, $I_{\mathrm{c},g}$ is I_{c} calculated over group g, and $\eta_i = \frac{\mathrm{b}_i x_i}{\mu^*}$ is the (revised) relative income of the i-th individual and $\eta_g = \frac{\mathrm{b}_g \mu_g}{\mu^*}$ that of the g-th group. $I_{\mathrm{c},g}$ is taken from equation 3, with g substituted for n, and $w_{\mathrm{c},g}$ from equation 5: which is to say, they are Shorrocks' indices, but calculated over the groups with revised mean incomes and their corresponding weights.

[18]Note that, possibly perversely, increasing the weight on a wealthy group is the way to be more averse to inequality as it appears as if they have even more income (or derive more welfare from a given level of income).

[19]For $\mathrm{c} \neq 0$, the results will be a loss of the ideal properties in disaggregating.

5 An example from Brazil: labour-market-earnings inequality across white men and non-white women

This article's contributions are illustrated through the decline of income inequality that occurred in Brazil between 2003 and 2015, without inequality changes being fully explored. Focusing on the labour-market earnings of white men and non-white women, we explore the levels and the changes in measures of inequality that obtain for a range of inequality-aversion parameters (c) and inter-group comparison parameters (b). Three points emerge. First is that in measuring inequality there is a trade-off between interpersonal comparability, through the choice of b, and the intensity of preference for inequality-aversion, through c, as discussed in sections 2 and 3. This allows for the novel construction of "iso-inequality" curves as a function of b and c. Second, and as an immediate consequence of the first point, is that both the level of overall inequality and its decomposition are affected by the choice of parameters b and c. Third is that the measured *changes* in inequality, both in absolute and relative terms, strongly depend on b and c, which allows for the construction of iso-variation curves as a function of both parameters. It is shown, for example, that for certain (admittedly extreme) values of b and c, income inequality across white men and non-white women did not decrease during the studied period, or did so at most marginally. This highlights that if there are reasons to treat groups differently, which is at least implicit whenever a between-groups decomposition is used, results crucially depend on assumptions about inter-group comparability.

5.1 Income inequality in Brazil: an overview

Brazil has historically been, and continues to be, one of the most unequal countries in the world. The country's inequalities encompass several dimensions, with income and wealth as central axes but spanning issues such as access to public services, health outcomes and treatment by the police (Hone et al. 2017, Loureiro 2020a, Willis 2015). These inequalities, across economic and other dimensions, have clear ethnoracial and gendered aspects, which intertwine with the country's class structure (Alves 2018, Lovell 2006, Moraes Silva and Paixão 2014, Rezende and Lima 2004).

Household per capita income inequality experienced a period of decline in Brazil, from the late 1990s until 2015, even if it continues to be very high by international standards and has risen since then (Hoffmann and Oliveira 2014, Prates and Barbosa 2020). This process has been widely studied and some of its major characteristics have been established. Overall, the decrease of per capita income inequality was largely driven by the labour market, with the public pension system and redistributive policies also playing a role (Loureiro 2020a). In terms of redistributive policies, the conditional cash transfer programme Family Allowance (*Programa Bolsa Família*) raised incomes at the very bottom of the distribution, whilst the Continuous Welfare Benefit (*Benefício de Prestação Continuada*), paid to the elderly poor, provided income support closer to the median income. Greater pension coverage, and real gains in low-value pensions, point to the progressive role of pensions.

The main driver of falling inequality, however, was an equalisation of labour-market earnings (Hoffmann and Oliveira 2014). The leading direct policy instrument in this process has been the minimum wage, which rose by 83% between January 2003 and December 2015 (in real terms, deflated by the National Index of Consumer Prices – INPC). This increased the wages of low-paid formal employees and, through a 'lighthouse effect', also raised earnings around the first quartile to the median (Maurizio and Vázquez 2016). Rising minimum wages have also been a determinant of the distribution of public pensions as they index the floor of pension benefits.

The other two main processes that contributed to falling labour-market inequality were a result of Brazil's overall pattern of growth during this period: the formalisation of, and real wage gains for, low-skilled occupations (Loureiro 2020b, Rugitsky 2019). With the initial increase of incomes in the bottom of the distribution, through greater government transfers and rising minimum wages, the demand for wage-goods and -services (i.e. goods and services typically consumed by relatively low-paid workers) rose. As these goods are mostly produced domestically in Brazil, this increased the demand in the corresponding occupations necessary to produce them, largely low-skilled. Heating this segment of the labour market fuelled real wage gains and labour formalisation for low-skilled workers – in turn re-igniting the demand for wage-goods and thus adding steam to the growth and redistribution cycle; in short, a cumulative causation process within the low-waged and poor.

Through these processes, income differentials between groups defined by their gender and racial self-identification also fell (Lima et al. 2013, Lima and Prates 2019). One key reason is that women, the black and brown population, and particularly black and brown women, are more likely to be in precarious positions in the labour market.[20]. As such, these underprivileged groups (compared to white men) are more likely to be paid close to the minimum wage, to be in informal positions, and to be in low-skilled occupations producing wage-goods (these occupations are gendered female and racialised black and brown in Brazil). There thus was a decrease of gender-race income gaps between 2003 and 2015, as a rising real minimum wage, low-skilled labour formalisation, and a heated labour market in wage-goods sectors all disproportionately benefited underprivileged groups, reducing labour-market inequality.

However, to what extent are conclusions over the implications for measurement of inequality dependent on the implicit assumption that one unit of income for, say, a white man, represents the same welfare that it does for a black woman? We investigate this in the following sub-section, showing that under most assumptions the broad direction of the inequality-reduction process is maintained, but that important aspects of it vary substantially depending on the degree of inter-personal comparability and inequality-aversion.

[20] For a discussion of the significance of skin colour and racial categories in Brazil, see Telles (2004). Joining the two survey categories – black and brown – is a standard procedure in statistical analyses because their economic characteristics are very similar. Additionally, there have been growing calls from social movements to build a common identity that encompasses black and brown individuals. In light of this, the article uses non-white as shorthand for black and brown.

5.2 Analysis

Data from the National Household Sample Survey (*Pesquisa Nacional por Amostra de Domicílios* – PNAD) are used to explore the inequality of labour-market earnings in Brazil. As discussed in sub-section 5.1, changes in the labour market were the key driver of falling inequality during the studied period, which justifies this empirical focus. Labour-market earnings in the PNAD include not only wages and salaries, but also other forms of income associated with an occupation, such as the earnings of employers and the self-employed (wages are used interchangeably as a shorthand for labour-market income).[21] Only the respondents' main occupation is taken into account in the empirical estimations, and only strictly positive wages are included.

The population is split into two mutually exclusive groups, white men and non-white women (see footnote 20). Sex (male or female) and race are based on respondent self-identification. Those who do not fall into these two categories (e.g. white women or indigenous individuals of any sex) are excluded from the analysis, so, in effect, this sub-sample of white men and non-white women is treated as the whole sample.

Table 1 reports descriptive statistics for the two periods. Incomes grew between 2003 and 2015 across the whole distribution, for both groups and for the overall population. Bottom incomes grew much more strongly than middle or top incomes, indicating that the decrease of inequality was driven comparatively more by an attenuation of deprivation than by a decrease of privilege (in absolute terms, for example, the 95th percentile grew more than the 75th or any other below it). Regarding the two groups, the mean income of non-white women grew by 66.3%, more than twice the 31.4% of white men, and the increase was more spread throughout the distribution (e.g. the group's 95th percentile increased by 52.0%, compared to 15.8% for white men). This made the relative income of white men compared to that of non-white women decrease from 2.9 to 2.3 during the period.

It can also be seen that other socioeconomic characteristics of white men and non-white women differed substantially. The latter were much more likely to be in precarious positions in the labour market; in 2015, they were about twice as likely as white men to be informal employees (27.6% against 14.0%) and nearly five times less likely to be employers (1.5% against 7.1%). Furthermore, non-white women on average spent nearly 19 hours per week performing unpaid domestic labour, compared to about five for white men, highlighting other forms of non-monetary inequality that affected the two groups. These dimensions, such as the burden of unpaid domestic labour and the relative chances of being informal employment, decreased much less than income differentials during the studied period. This serves in part as motivation to consider a range of inter-group comparison parameters (b) to measure the income inequality within and between groups – to reflect, for example, the greater burdens attached to non-white women – although the list of factors provided here is merely illustrative and could be greatly expanded.

Restricting for the moment b to the same value for both groups, Table 2 shows the

[21] Similarly to other household surveys, the PNAD under-represents top incomes in general, and capital-based incomes in particular, which leads to downward-biased estimates of inequality (for a comparison with tax-based data, see Medeiros and Castro 2018).

Table 1: Descriptive statistics of labour-market earnings in Brazil for white men and non-white women, 2003 and 2015

	2003			2015			$\Delta_{2015-2003}$		
	WM	NWW	Total	WM	NWW	Total	WM	NWW	Total
Observations	44,462	28,494	72,956	36,633	34,355	70,988	-7,829	5,861	-1,968
Population %	65.5	34.5	100.0	55.4	44.6	100.0	-10.1	10.1	0.0
Employers	8.0	1.4	5.7	7.1	1.5	4.6	-0.9	0.1	-1.1
Self-employed	27.3	20.9	25.1	26.8	20.0	23.8	-0.5	-0.9	-1.4
Formal employees	45.6	40.2	43.7	52.1	50.9	51.6	6.5	10.7	7.9
Informal employees	19.1	37.5	25.4	14.0	27.6	20.1	-5.1	-9.9	-5.4
Weekly hours of household labour*	4.5	20.4	9.9	5.1	18.8	11.2	0.7	-1.6	1.3
Income mean	1,967	686	1,525	2,585	1,141	1,941	31.4%	66.3%	27.2%
Income perc. 25	592	296	474	1,000	600	788	68.8%	102.6%	66.3%
Income perc. 50	987	474	790	1,500	840	1,200	52.0%	77.3%	52.0%
Income perc. 75	1,974	770	1,579	2,700	1,200	2,000	36.8%	55.8%	26.6%
Income perc. 95	6,910	1,974	5,528	8,000	3,000	6,000	15.8%	52.0%	8.5%
Relative income, overall population	1.29	0.45	1.00	1.33	0.59	1.00	0.04	0.14	0.00
Relative income, other group**	2.87	0.35	–	2.27	0.44	–	-0.60	0.09	–
Income %	84.5	15.5	100.0	73.8	26.2	100.0	-10.7	10.7	0.0

Source: Prepared by the authors based on data from PNAD-IBGE (2003/2015).
Notes: *This refers to unpaid domestic labour carried out in the individual's household.
**Relative income with regards to the overall population is the mean income of the group in question divided by overall mean income, $\frac{\mu_g}{\mu}$, whilst relative to the other group is $\frac{\mu_g}{\mu_{i \neq g}}$.
Incomes deflated by INPC, 2015 values. WM refers to white men, NWW to non-white women.

levels and changes of inequality for different inequality-aversion parameters, c. Column 1 shows overall measure, the sum of Columns 2 and 3 which are total within and between group inequality. Weighted sums of within group inequality (Columns 4 and 5) yield Column 2. Given that low-wage gains drove the decrease of inequality, the index falls more for lower values of c; 33.5% for c = −1, for example, compared to 12.0% for c = 2. This decrease was dominated, in absolute terms, by the much larger falls of within-group inequalities, although the within- and between-groups components decreased by similar amounts relative to their initial values. As a result, the share of inequality accounted for by each of the within and between components was approximately the same at the beginning and end of the period (measured by $\frac{I^W}{I}$ and $\frac{I^B}{I}$), ranging for between groups from a maximum of about 17% for c = 0, which is more sensitive to change in the middle of the distribution, to about 5% for c = 2 and 10% for c = −1. Whilst the relatively low share of between-groups inequality might come as a surprise, given the relative income of white men compared to non-white women is 2.27 with similar population-shares (in

Table 2: Labour-market-earnings inequality with different inequality-aversion parameters (c) in Brazil, 2003 and 2015

c	2003					2015					$\%\Delta_{2015-2003}$				
	I	I^W	I^B	I_{whm}	I_{nww}	I	I^W	I^B	I_{whm}	I_{nww}	I	I^W	I^B	I_{whm}	I_{nww}
-1.00	1.29	1.16	0.14	0.96	0.87	0.86	0.77	0.09	0.67	0.65	-33.5	-33.1	-36.4	-29.8	-25.5
-0.50	0.78	0.66	0.12	0.65	0.55	0.56	0.47	0.08	0.48	0.41	-28.6	-27.9	-32.2	-25.6	-24.3
0.00	0.62	0.51	0.11	0.55	0.44	0.46	0.39	0.08	0.43	0.33	-24.6	-23.8	-28.1	-21.4	-23.3
0.50	0.59	0.49	0.10	0.53	0.41	0.46	0.39	0.07	0.44	0.32	-21.9	-21.4	-24.3	-18.3	-23.2
1.00	0.67	0.58	0.09	0.60	0.46	0.54	0.47	0.07	0.51	0.34	-20.0	-19.9	-20.6	-15.8	-24.8
1.50	0.91	0.83	0.08	0.80	0.60	0.75	0.68	0.07	0.70	0.43	-17.8	-17.8	-17.2	-12.5	-28.8
2.00	1.55	1.47	0.08	1.28	0.98	1.36	1.29	0.07	1.22	0.62	-12.0	-11.9	-14.1	-5.2	-36.1

Source: Prepared by the authors based on data from PNAD-IBGE (2003/2015).
Note: Inter-group comparison parameter (b) set to equal weights for both groups.

2015), differences in relative income within groups were even starker. In the case of white men, the top half earned on average 4.26 times that the bottom half in 2015, while the corresponding value for non-white women was 3.12 (see also the percentiles in table 1).

It should also be noticed that the distribution of earnings for non-white women always remain more equal than for white men, especially for high values of c (i.e. top-end inequality is comparatively higher amongst white men). Low-earning individuals were present in both groups, leading to comparable inequality indices for c $= -1$, particularly after the gains in the bottom of distribution that occurred up to 2015. In this vein, whilst $I_{-1,whm} = 0.67$ and $I_{-1,nww} = 0.65$ in 2015, the absolute and relative gap between the indices for the two groups rose steadily for higher values of c, reaching nearly double the value of each other for c $= 2 : I_{2,whm} = 1.22$ and $I_{2,nww} = 0.62$.

Table 3 reports inequality indices for a range of inter-group comparison parameters, fixing the inequality-aversion parameter c $= 0$. As discussed in section 4, the latter leads to unambiguous interpretations of the decomposition, as the weights of the within-groups component are population-shares and, as such, independent of the groups' relative incomes (which then determine the between-groups inequality component). In light of this, varying b for c $= 0$ does not affect I^W but only changes overall inequality through its impact on I^B. This is because within-group income is only rescaled and so inequality remains unchanged by changes in b. That is why the column values for I^W remain constant, 0.51 for 2003 and decreasing substantially to 0.39 for 2015 (with both groups contributing substantially to this inequality if slightly more from white males, see Table 2). Further, in Table 3, values of b lower than 0.5 (perversely favouring white men over non-white women) are included for the sake of illustration, but do not carry great practical significance as they would reflect a view that white men are, as a group, under-privileged compared to non-white women (leading, for example, to a null I^B when b is 0.30, close to the inverse of the two group's relative incomes). In short, the point is to show how inequality increases as we weight women's income less in the measure (b higher), as if their income were lower.

Even small variations of b lead to substantial reassessments of inequality. For b $= 0.6$,

Table 3: Labour-market-earnings inequality with different inter-group comparison parameters (b) in Brazil, 2003 and 2015

	2003			2015			$\%\Delta_{2015-2003}$		
b	I	I^W	I^B	I	I^W	I^B	I	I^W	I^B
0.10	0.67	0.51	0.16	0.62	0.39	0.23	-7.5	-23.8	44.7
0.20	0.52	0.51	0.01	0.43	0.39	0.04	-18.0	-23.8	210.8
0.30	0.51	0.51	0.00	0.39	0.39	0.00	-24.5	-23.8	-97.7
0.40	0.55	0.51	0.04	0.41	0.39	0.02	-26.1	-23.8	-53.0
0.50	0.62	0.51	0.11	0.46	0.39	0.08	-24.6	-23.8	-28.1
0.60	0.70	0.51	0.20	0.55	0.39	0.17	-21.2	-23.8	-14.3
0.70	0.82	0.51	0.31	0.68	0.39	0.29	-16.6	-23.8	-4.6
0.80	0.97	0.51	0.46	0.86	0.39	0.48	-10.9	-23.8	3.2
0.90	1.23	0.51	0.72	1.18	0.39	0.79	-3.7	-23.8	10.5

Source: Prepared by the authors based on data from PNAD-IBGE (2015).
Note: Inequality-aversion parameter (c) set to 0. b multiplies the income of white men, and (1-b) that of non-white women.

meaning white men enjoy 50% more welfare for a given level of income than non-white women (60% relative to 40%, respectively), I^B roughly doubles in each year (from 0.08 to 0.17 in 2015). For b = 0.7, the between-groups component represents about 40% of total inequality. Across the two years, it straddles contributing 50% of inequality for b = 0.8. Although b = 0.8 represents what might be considered an extreme value that would be hard to justify, indicating non-white women would require four times as much income to obtain a given level of welfare (80% relative to 20%), it does lead to analytically interesting results.

In comparing change between the two years, it is important to bear two effects in mind. On the one hand, relative income per capita between the two groups ($\frac{\mu_{whm}}{\mu_{nww}}$) has decreased due especially to the income-increasing factors already discussed. On the other hand, there is an increase in the numbers (i.e. a higher population-share) on this increased but still lower-income group, tending to increase inequality by adding more individuals on lower incomes ($p_{nww} = 34.5\%$ in 2003 and 44.6% in 2015).[22] By the same token, changing b will have two shifting effects on inequality measures between the two dates, reflecting both increases in relative incomes for the lower paid and greater numbers of them. As a result, whilst overall inequality reduces for each value of b between the two years, the second effect on I^B outweighs the first once b reaches 0.8 (increasing between-groups

[22]Notice that the terms of the summation of I^B are the mean income of the group relative to that of the whole population, $\eta_g = \frac{\mathrm{b}_g \mu_g}{\mu^*}$, and not relative to each other, $\frac{\mathrm{b}_g \mu_g}{\mathrm{b}_h \mu_h}$. With this, even for b = 0.5, in light of the entry of non-white women into the paid labour market it can be seen that η_{whm} rose from 1.29 to 1.33 between 2003 and 2015 whilst η_{nww} also rose, from 0.45 to 0.59, although $\frac{\mu_{whm}}{\mu_{nww}}$ declined from 2.87 to 2.27 (see table 1).

inequality with I^B moving from 0.46 to 0.48), and even more so when b=0.9 (from 0.72 to 0.79).

More generally, there is a neat way of bringing together the analysis of combined variations in b and c together. Figure 1 shows 'iso-inequality' curves for 2003 and 2015, constructed by calculating the inequality indices that obtain for the ranges b = $[0.1, 0.9]$ and c = $[-1.0, 2.0]$ and drawing curves that connect the points for which I equals selected values. What this shows is exactly how the degrees of inequality-aversion and inter-group comparability can be traded off against one another for a given level of inequality. For example, in 2015, the inequality index is the same if b = 0.5 (so that income is not revised) and c = -0.62 and if b = 0.67 and c = 0.59. Increasing aversion to bottom-end inequality, by reducing c from 0.59 to -0.62, has therefore the same but opposite effect on the measure as revising the income of white men to be approximately double that of non-white women. It can also be seen that there is a higher density of curves, so that inequality is more sensitive to the parameters, with low values of c and high values of b – i.e. when the index is sensitive to bottom-end inequality and non-white women are positioned in the extreme bottom of the distribution of revised income. Finally, the index is locally inelastic with respect to b or to c over certain ranges, when the inclination of the curves is parallel to a certain axis. In 2003, for example, I is insensitive to c if b $\approx$ 0.8 and c $\approx$ 0.75. In 2015, I is insensitive to b if b $\approx$ 0.31 and c $\approx$ 0, or if b $\approx$ 0.25 and c $\approx$ 1.5.[23]

Such considerations are more fully covered in Figure 2. It shows the relative changes in inequality between the two periods for different values of b and c, where the relative change is defined as $\Delta^{\%}_{\mathrm{b,c}} = 100 * \left(\frac{I_{2015,\mathrm{b,c}} - I_{2003,\mathrm{b,c}}}{I_{2003,\mathrm{b,c}}}\right)$. Working with relative changes has the advantage of presenting easily interpreted cardinal values (i.e. how much did inequality decrease relative to its starting point), which is not usually the case with I_{c}, whose cardinal interpretation is not straightforward.[24] With this, a clearer assessment is possible of how decisions about inequality-aversion and interpersonal comparability affect the trajectory of inequality in Brazil.

Restricting ourselves to the range of b $\geq$ 0.5, inequality fell proportionally more for lower values of c. This is because, as has been mentioned before, the decrease of inequality was driven by gains in the lower end of the distribution, with top-end inequality proving more resilient. Also, inequality fell less for higher values of b, whichever the value of c, although the sensitivity of the index depends on the combination of both parameters. Furthermore, although inequality decreased for most of the range of b and c under consideration, there are some exceptions. If we are particularly concerned with top-end inequality and have reason to consider that white men enjoyed a large multiple of monetary income compared to non-white women, overall inequality increased during the

[23]The value of b which minimises I is not always that which zeroes I^B (although this is necessarily the case for c = 0). This is so because, for c $\neq$ 0, an exponent of the groups' relative income enters into the weights of the summation of the *within*-groups component (see expression 5). With this, given that $I_{whm} > I_{nww}$ and $\eta_{whm} > \eta_{nww}$, driving b for a range below the point where the revised mean income of the groups is the same (so that $I^B = 0$) reduces overall inequality for high levels of c.

[24]This is the case, amongst others reasons, because I_{c} has a varying upper limit depending on c and is unbounded from above for c $\leq$ 0, even if its lower limit is always 0.

Figure 1: Iso-inequality curves for different inequality-aversion (c) and inter-group comparison (b) parameters, labour-market-earnings inequality in Brazil, 2003 and 2015

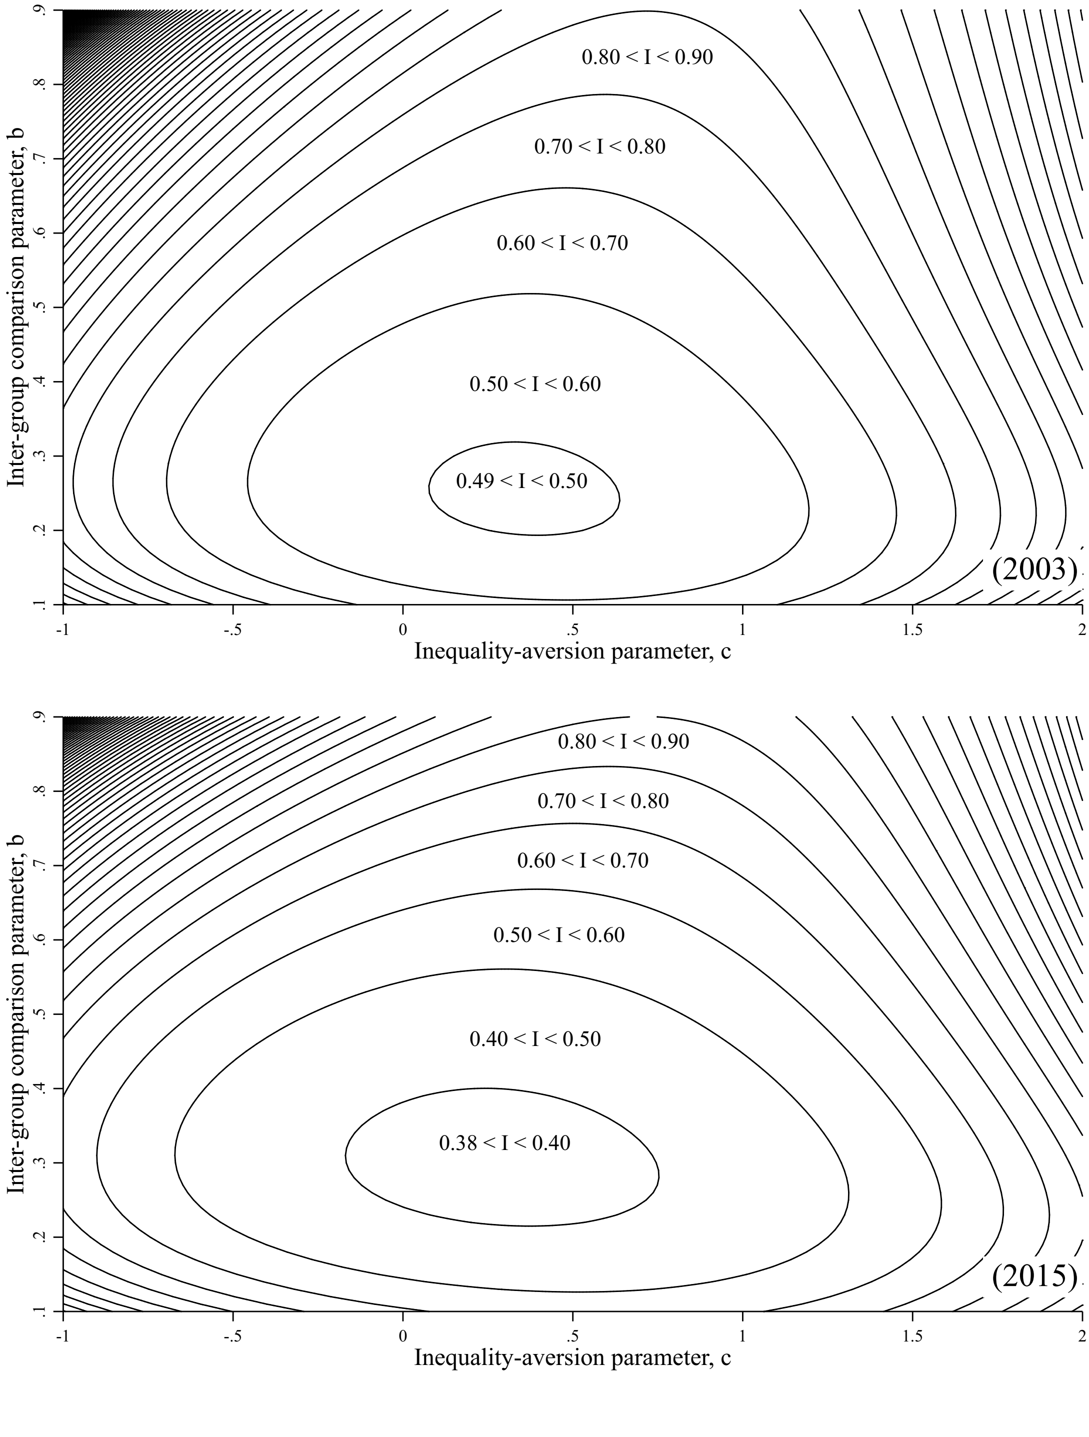

Source: Prepared by the authors based on data from PNAD-IBGE (2003/2015).
Note: b multiplies the income of white men, and (1-b) that of non-white women.

Figure 2: Iso-variation curves for percentage change in labour-market-earnings inequality in Brazil, 2003-2015, for different inequality-aversion (c) and inter-group comparison (b) parameters

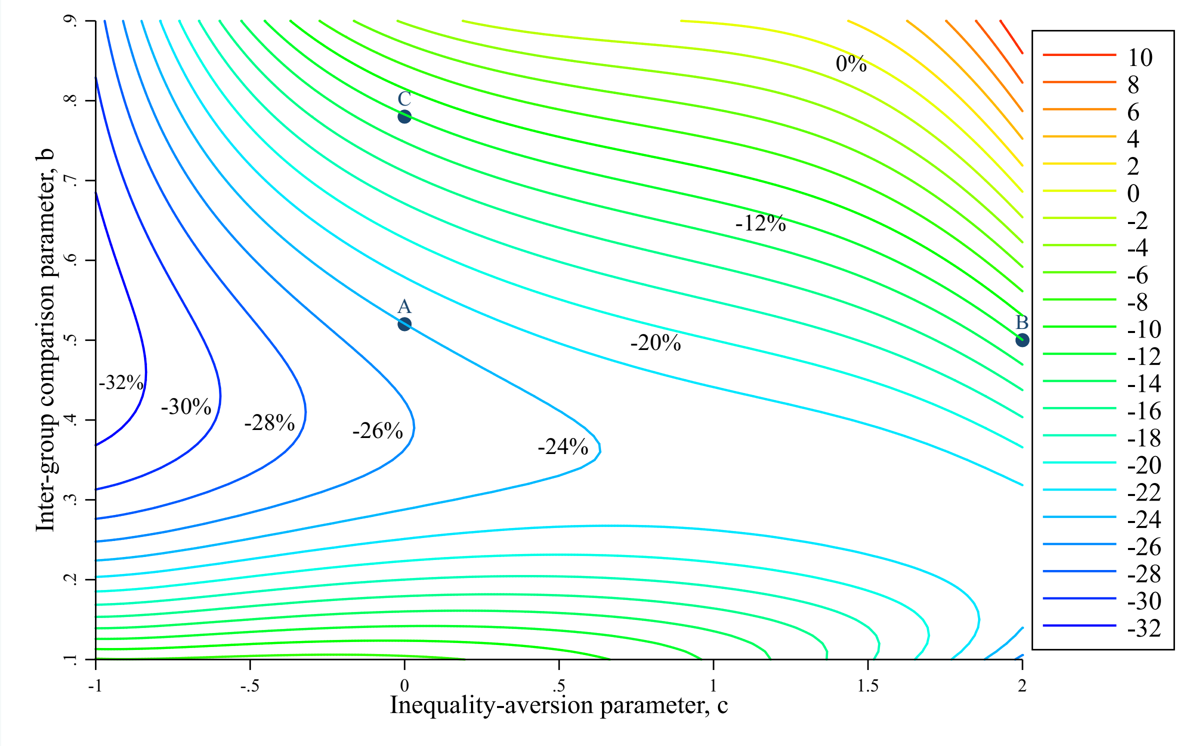

Source: Prepared by the authors based on data from PNAD-IBGE (2003/2015).
Note: b multiplies the income of white men, and (1-b) that of non-white women. Percentage change defined as $\Delta^{\%}_{\mathrm{b,c}} = 100 * \left(\frac{I_{2015,\mathrm{b,c}} - I_{2003,\mathrm{b,c}}}{I_{2003,\mathrm{b,c}}}\right)$.

period. This happens, for instance, if b > 0.85 and c > 1.5, or if b > 0.67 and c > 2.0,[25] although these are extreme values.

For illustrative purposes, follow two of the iso-variation curves, for $\Delta^{\%}_{\mathrm{b,c}} = -12$ and -24. We can see certain combination of parameters that produce the same assessment of how inequality varied across the two dates. For example, $\Delta^{\%}_{\mathrm{b,c}} = -12 \approx \Delta^{\%}_{0.5,2} \approx \Delta^{\%}_{0.67,1} \approx \Delta^{\%}_{0.8,0} \approx \Delta^{\%}_{0.98,-0.63}$. This shows how quickly ‘extreme’ values of b are required to maintain $\Delta^{\%}_{\mathrm{b,c}}$ constant as we alter the sensitivity to inequality: a transformation of the order of 50 to 1 in relative incomes is necessary to offset the impact of moving from c = 2 to -0.63. Alternatively, this also shows how strong is the impact of varying c, however much a standard procedure it might be in the literature, seen through its counterpart of varying inter-group comparability. Similarly, $\Delta^{\%}_{\mathrm{b,c}} = -24 \approx \Delta^{\%}_{0.5,0} \approx \Delta^{\%}_{0.67,-0.5} \approx \Delta^{\%}_{0.8,-0.75} \approx$

[25]For b = 0.5, there is Lorenz dominance between the 2015 and 2003 curves, but this ceases to be the case for higher levels of b.

$\Delta^{\%}_{0.98,-0.9}$.

Points $\mathrm{A}_{\mathrm{b}=0.5,\mathrm{c}=0}$, $\mathrm{B}_{0.5,2}$ and $\mathrm{C}_{0.8,0}$ in Figure 2 help to connect the analysis of these two iso-variation curves. We can see that moving to point A from B or from C has the same impact on the variation of inequality, i.e. $\left(\Delta^{\%}_{0.5,2} - \Delta^{\%}_{0.5,0}\right) \approx \left(\Delta^{\%}_{0.8,0} - \Delta^{\%}_{0.5,0}\right)$. In effect, moving from an inequality-aversion parameter of 2 to 0 (with b = 0.5) has the same impact on the proportional changes of inequality (i.e. a variation of 12 pp) as moving from an interpersonal comparability parameter of 0.5 to 0.8 (with c = 0).

To make the examples more concrete, let us consider the following three representative individuals. Working with untransformed incomes, in 2015, a non-white woman employed as an informal domestic worker had an average income of R\$ 585 ($\eta_{dom} = 0.301$), which positioned her at 0.145 of the ranked distribution of income. Someone precisely at the median had an income of R\$ 1,200 ($\eta_{med} = 0.618$, close to the mean income of a non-white female secretary or a white male warehouse operative), whilst a white man who employed ten or more workers had an average income of R\$ 11,331 ($\eta_{emp} = 5.838$), putting him at 0.985 of the ranked distribution.

Making use of the additive property of the inequality index, it is possible to calculate the relative impact of transfers between these individuals. If $T_{i\to j}$ denotes a small transfer from the i-th to the j-th individual, and making use of the additive property of the inequality index, we can explore how sensitive is I_c to transfers from the top to the middle, and from the middle to the bottom, of the distribution. With these values indicated above, the relative sensitivity of I_c to transfers from a white male employer to someone in the median, and from someone in the median to a non-white female informal domestic worker is:

$$
\begin{aligned}
&\mathrm{c} = -1, && T_{emp\to med} = 0.308\ T_{med\to dom}\\
&\mathrm{c} = 0, && T_{emp\to med} = 0.850\ T_{med\to dom}\\
&\mathrm{c} = 1, && T_{emp\to med} = 3.125\ T_{med\to dom}\\
&\mathrm{c} = 2, && T_{emp\to med} = 16.473\ T_{med\to dom}
\end{aligned}
$$

The magnitudes of these results are striking, but they illustrate what is at stake by choosing different forms of inequality-aversion. A greater concern with top-end inequality, associated with higher values of c, is justified if one is interested in, say, the impact of income inequality on the distribution of political influence – arguably a phenomenon related to very high levels of relative income (and wealth). Nevertheless, if c = 2, the index puts a very low penalty indeed on relative poverty, just a sixteenth of the impact of moving income at the lower end. Alternatively, a greater concern with poverty aversion, associated with low values of c, comes at the expense of a lower penalty for top-end inequality, arguably excessively low if c = −1.

Returning to the example $\left(\Delta^{\%}_{0.5,2} - \Delta^{\%}_{0.5,0}\right) \approx \left(\Delta^{\%}_{0.8,0} - \Delta^{\%}_{0.5,0}\right)$, we can now flesh out these results. This indicates that changing our aversion to inequality from c = 2 (in which transfers from a white male employer to someone in the median should be valued at about 16 times more than from the median to an informal non-white domestic worker) to c = 0 (in which these same transfers would be valued at 0.85 each other), the movement from point B to A, has the same impact on the variation of inequality as considering that the relative income of our two groups should be divided by four (from b = 0.8 to 0.5), moving from C to A. Whilst these higher ranges of b and c might be unjustified in most circumstances, the variations considered are nevertheless equivalent to each other in terms of their impact on the relative change of inequality that occurred, highlighting once more the practical consequences of the equivalences between inequality-aversion and interpersonal comparability. Restricting ourselves to smaller ranges, with $0.5 \leq b \leq 0.67$ and $-0.5 \leq c \leq 1.0$, overall inequality is seen to have varied between -28.6% and -12.2%.

6 Final remarks

This article has provided a link between social choice theory and the measurement of inequality, fields which share an intrinsic implicit connection that has, nevertheless, been little explored in the literature. It was shown how, when ranking alternatives in the context of social choice, that there are two routes out of extreme restrictions of the Pareto principle: varying the intensity of preference for each individual and varying the degree of interpersonal comparability. It was then shown how these two paths map onto the measurement of inequality as, respectively, the degree of inequality-aversion and the interpersonal comparability of incomes (or welfare).

Over thirty years ago, Fine (1985) pointed out that the Atkinson measure of inequality without symmetry led to results as if there were linear interpersonal comparability parameters (denoted by b here). This result has mainly been unobserved. A decade later, if drawing on his doctoral thesis, Fine (1974), Fine (1996) pointed to the equivalence between interpersonal comparability (b) and inequality-aversion (denoted by c here) – also mainly unobserved. Revisiting these issues, Fine and Loureiro (2020) have forged a reconciliation, a duality, between between the Atkinson and Shorrocks approaches to measuring inequality, offering novel insights around their respective measures and the conditions that yield them. What this paper has done is brought these three sets of results together to create a powerful framework for measuring inequality in the context of its decomposition across different groups. The Shorrocks restriction on requiring the inequality-aversion parameter c to be zero (to allow unambiguous decomposition) has been compensated for by varying the interpersonal comparability parameter, b. Technically, then, this article has proposed a measure of inequality that drops the axiom of symmetry whilst allowing for between-groups additive decomposability. As suggested in Section 2, possibly the negative ethos of impossibility attached to social choice theory has delayed its integration with the statistical measurement approach to inequality.

This article then illustrated potential results of such developments through an application to Brazil. It investigated the widely-documented decrease of labour-market income inequality that occurred in the country between 2003 and 2015, restricting the sample to white men and non-white women and decomposing inequality into these two groups' contributions. A range of inequality-aversion and inter-group comparability parameters were explored, considering that white men enjoy a multiple of their monetary income when compared to non-white women. This was justified in light of the disadvantages non-white women experience that are not captured in their income (e.g. higher labour informality, lower prospects for job progression, a greater burden of unpaid domestic labour). Based on this, it was shown that not only does the level of inequality at a given point depend on b, but also, less anticipated, so does its variation. The equivalence of varying b and c was also demonstrated, for example by showing that moving from an inequality-aversion parameter of c = 2 to 0 (whilst keeping b equal for both groups) has the same effect on the proportional change of inequality as moving from an equal comparability parameter to multiplying the relative income of white men by a factor of 4 (whilst keeping c = 0).

These results do not question the overall trajectory of inequality in Brazil between 2003 and 2015, but they do offer important qualifications. Whilst it is only for extreme values of b and c that income inequality is seen to have increased, the intensity of the fall does change substantially within a reasonable range of parameters. Considering that white men enjoy between the same and twice as much welfare from a given level of income, and setting inequality-aversion to $-0.5 \leq c \leq 1.0$, overall inequality is seen to have varied between -28.6% and -12.2% over the period, a substantial range by any account. It highlights that our general analysis of the interpersonal comparability and inequality-aversion trade-off needs to be addressed explicitly in any empirical measurements of, or judgements over, how much inequality has changed.

7 References

Alves, J. A., 2018. *The Anti-Black City: Police Terror and Black Urban Life in Brazil.* Minneapolis: University of Minnesota Press.

Arrow, K. 1951. *Social Choice and Individual Values.* New Haven: Yale University Press.

Atkinson, A. 1970. On the Measurement of Inequality. *Journal of Economic Theory*, 2(3), 244-263.

Backhouse, R., Baujard, A. and Nishizawa, T. 2020. Revisiting the History of Welfare Economics. *HAL* Id: halshs-02937994, https://halshs.archives-ouvertes.fr/halshs-02937994

Bossert, W., Suzumura, K. 2020. Positionalist voting rules: a general definition and axiomatic characterizations. *Social Choice and Welfare*, 55, 85–116.

Cato, S. 2020. Extending the Intersection Approach. *Journal of Human Development and Capabilities*, 21(3), 230-248.

Deutsch, J. and Silber, J., 1999. Inequality Decomposition by Population Subgroups and the Analysis of Interdistributional Inequality. *In:* Silber, J. ed. *Handbook of Income Inequality Measurement.* Boston: Springer Science, 363-404.

Fine, B. 1974. *Individual Decisions and Social Choice*, unpublished Phd thesis, University of London.

Fine, B. 1985. A Note on the Measurement of Inequality and Interpersonal Comparability, *Social Choice and Welfare*, 1(4), 273-77.

Fine, B. 1996. Reconciling Interpersonal Comparability and the Intensity of Preference for the Utility Sum Rule. *Social Choice and Welfare*, 13(2), 319-25.

Fine, B. 2004. Economics and Ethics: Amartya Sen as Point of Departure. *The New School Economic Review*, 1(1), 151-62.

Fine, B. 2016. *Microeconomics: A Critical Companion*. London: Pluto.

Fine, B., 2018. Collective Choice and Social Welfare: Economics Imperialism in Action and Inaction. *Ethics and Social Welfare*, 12(4), 393-399.

Fine, B. and Fine, K. 1974a. Social Choice and Individual Ranking, Part I. *Review of Economic Studies*, 41(3), 303-322.

Fine, B. and Fine, K. 1974a. Social Choice and Individual Ranking, Part II. *Review of Economic Studies*, 41(4), 459-475.

Fine, B. and Loureiro, P. M. 2020. A Note on the Relationship between Additive Separability and Decomposability in Measuring Income Inequality. *Review of Social Economy.* online first, https://www.tandfonline.com/doi/full/10.1080/00346764.2020.1802055

Fine, B. and D. Milonakis. 2009. *From Economics Imperialism to Freakonomics: The Shifting Boundaries between Economics and Other Social Sciences.* London: Routledge.

Hoffmann, R. and Oliveira, R. B. 2014. The Evolution of Income Distribution in Brazil in the Agricultural and the Non-agricultural Sectors. *World Journal of Agricultural Research,* 2(5), 192-204.

Hone, T.*, et al.* 2017. Association Between Expansion of Primary Healthcare and Racial Inequalities in Mortality Amenable to Primary Care in Brazil: A national Longitudinal Analysis. *PLOS Medicine,* 14(5), e1002306.

Lima, M. and Prates, I., 2019. Racial Inequalities in Brazil: A Persistent Challenge. *In:* Arretche, M. ed. *Paths of inequality in Brazil: a half-century of changes.* Cham: Springer, 113-134.

Lima, M., Rios, F. and França, D., 2013. Articulando Gênero e Raça: a Participação das Mulheres Negras no Mercado de Trabalho (1995-2009). *In:* Marcondes, M. M.*, et al.* eds. *Dossiê Mulheres Negras: Retrato das Condições de Vida das Mulheres Negras no Brasil.* Brasília: IPEA, 53-80.

Loureiro, P. M. 2020a. Class Inequality and Capital Accumulation in Brazil, 1992-2013. *Cambridge Journal of Economics,* 44(1), 181–206.

Loureiro, P. M. 2020b. Social Structure and Distributive Policies in Brazil Under the PT Governments: a Poverty-Reducing Variety of Neoliberalism. *Latin American Perspectives,* 47(2), 65-83.

Lovell, P. A. 2006. Race, Gender, and Work in São Paulo, Brazil, 1960-2000. *Latin American Research Review,* 41, 63-87.

Maurizio, R. and Vázquez, G. 2016. Distribution Effects of the Minimum Wage in Four Latin American Countries: Argentina, Brazil, Chile and Uruguay. *International Labour Review,* 155(1), 97-131.

Medeiros, M. and Castro, F. Á. 2018. A Composição da Renda no Topo da Distribuição: Evolução no Brasil Entre 2006 e 2012, a Partir de Informações do Imposto de Renda. *Economia e Sociedade,* 27(2), 577-605.

Moraes Silva, G. and Paixão, M., 2014. Mixed and Unequal: New Perspectives on Brazilian Ethnoracial Relations. *In:* Telles, E. ed. *Pigmentocracies: Ethnicity, Race, and Color in Latin America.* Chapel Hill: University of North Carolina Press, 172-217.

Prates, I. and Barbosa, R. J. 2020. The Impact of COVID-19 in Brazil: Labour Market and Social Protection Responses. *The Indian Journal of Labour Economics.* online, https://doi.org/10.1007/s41027-020-00252-3.

Rezende, C. B. and Lima, M. 2004. Linking Gender, Class and Race in Brazil. *Social Identities,* 10(6), 757-773.

Rugitsky, F., 2019. Questão de Estilo: a Mudança Estrutural para a Igualdade e seus Desafios. *In:* Chiliatto-Leite, M. V. ed. *Alternativas para o Desenvolvimento Brasileiro: Novos Horizontes para a Nudança Estrutural com Igualdade.* Santiago: CEPAL, 75-95.

Sen, A. 1970a. *Collective Choice and Social Welfare.* San Francisco: Holden-Day.

Sen, A. 1970b. Interpersonal Aggregation and Partial Comparability. *Econometrica*, 38(3) 393-409.

Sen, A. 1973. *On Economic Inequality.* Oxford: Clarendon Press.

Sen, A. 1999. The Possibility of Social Choice, Nobel Lecture. *American Economic Review*, 89(3), 349-78.

Sen, A. 2017. *Collective Choice and Social Welfare.* London: Penguin, second edition.

Shorrocks, A. 1980. The Class of Additively Decomposable Inequality Measures. *Econo-*

metrica, 48(3), 613-625.

Telles, E., 2004. *Race in Another America: The Significance of Skin Color in Brazil.* Princeton: Princeton University Press.

Willis, G. D., 2015. *The Killing Consensus: Police, Organized Crime, and the Regulation of Life and Death in Urban Brazil.* Oakland: University of California Press.