SOME CONSIDERATIONS RELATING TO THE MEASUREMENT AND MITIGATION OF INEQUALITY

S. SUBRAMANIAN

WORKING PAPER N°2025/14

JULY 2025

WORLD INEQUALITY LAB

Title: Some Considerations Relating to the Measurement and Mitigation of Inequality

Author Details:
Name: S. Subramanian.
Institutional Affiliation: Independent Scholar; formerly, Madras Institute of Development Studies.
Address: 36-H, North Parade Road, St Thomas' Mount, Chennai – 600 016, Tamil Nadu, India.
e-mail address: ssubramanianecon@gmail.com

Abstract: Amongst the most widely used measures of income or wealth inequality is the Gini coefficient *G*. In recent times, there has been a rise in popularity of alternative measures such as the Palma Ratio (which is the ratio of the income-share of the top 10 per cent of a distribution to that of the bottom 40 per cent), or simply the income-shares of 'top incomes' (the top 1 per cent, or 0.1 per cent, or even 0.01 per cent), as in the work of Atkinson, Piketty and others. A particularly simple measure which reckons 'top incomes' is *R*, the proportionate shortfall of the mean income from the richest person's income. Measures such as *G*, which take stock of the distribution in its entirety, may be called 'across the board' indices, while measures such as *R*, which focus attention on the tail(s) of the distribution, may be called 'tailender' indices. Rather than see these categories of measures as being in opposition to each other, the present paper suggests that there is a case for combining them in a composite index. One such measure, which combines *G* and *R*, is the index *D*, given (for 'large' populations) by $D = R + (1\text{-}R)G$. The inequality index *D* is derived in this paper by way of an elementary extension of what is known as the Sen-Shorrocks-Thon poverty index into a well-defined index of inequality. It turns out that for distributions in which income is heavily concentrated at the upper end, the value of *D* is significantly influenced by the value of *R*, and that a natural approach to the reduction of inequality would be to cap top incomes. The paper suggests that considerations of both measurement and political morality would espouse the distributional doctrine of 'limitarianism', as proposed by Ingrid Robeyns, as an inevitable concomitant of the mitigation of inequality.

Key Words: 'across the board' measures; 'tailender' measures; Gini coefficient; proportionate shortfall of mean income from richest person's income; composite measure of inequality; capping top incomes; limitarianism

JEL Classification: D31, D63

Acknowledgements: This paper has benefited from earlier versions having been read and commented upon by John Creedy, Jayati Ghosh, Nicole Hassoun, Thomas Piketty and Anmol Somanchi, the last of whom, in addition, has been of considerable material help in the production of the paper. To these people, I offer grateful thanks, without seeking to implicate them in the paper's errors and shortcomings.

WIL working papers are circulated for discussion and comment purposes. Short sections of text may be quoted without explicit permission provided that full credit is given to the author(s). CC BY-NC-SA 4.0
How to cite this working paper: Subramanian, S., Some Considerations Relating to the Measurement and Mitigation of Inequality, World Inequality Lab Working Paper 2025/14

Some Considerations Relating to the Measurement and Mitigation of Inequality

S. Subramanian

1. Introduction

Amongst the most popular and widely-employed measures of income inequality to be found in the literature is the Gini coefficient. In recent times, other indicators of inequality, such as the Palma Ratio (Palma, 2011) and 'top income shares'—such as in the work, in particular, of economists like Alvaredo, Atkinson, Chancel, Piketty, Saez and Zucman[1]—have acquired considerable prominence. The Gini coefficient, like many other indices of inequality, is derived from a consideration of an income distribution in its entirety. By contrast, the Palma measure, which is the ratio of the income-share of the richest 10 per cent of a population to the income share of the poorest 40 per cent, concentrates attention on the tails of the distribution, while the 'top incomes' literature has a focus on the upper tail of the distribution, concerned as it is with the income shares and distributions of the top 1 per cent, the top 0.1 per cent, and even the top 0.01 per cent. In the Palma Ratio, the exclusion from consideration of the 'middle' fifth-to-ninth deciles of the distribution is rationalised by appeal to the observed relative stability of the income-share of these middle classes, at around 50 per cent—an empirical regularity that appears to hold across countries at a given point in time, and over time for a given country (Cobham, Schlogl and Sumner, 2015). The engagement of the 'top incomes' literature with the concentration of incomes and wealth with the very rich is motivated, in the most general sense, by the notion that "…people have a sense of fairness and care about the distribution of economic resources across individuals in society"; this is reflected in the view that "…all advanced economies have set in place redistributive policies such as taxation—and in particular progressive taxation, and transfer programmes, which effectively redistribute a significant share of National Product across income groups" (Atkinson, Piketty and Saez, 2011; p.7).

[1] For a smattering of studies in this area, the reader is referred to, among others, Atkinson, Piketty and Saez (2011), Alvaredo, Atkinson, Piketty and Saez (2013), Zucman (2019), Chancel and Piketty (2021), Bharti, Chancel, Piketty and Somanchi (2024), and Saez and Zucman (2024).

Commentators such as Cobham et al (2015) have found much merit in the Palma Ratio, while others, like Milanovic (2015), have displayed less enthusiasm for it. We do not here propose to delve in detail into the pros and cons of the Gini and Palma / 'top incomes' indices, nor even to review, with any degree of scrutiny, the virtues/deficiencies of these alternative approaches to the assessment of inequality. As noted earlier, the Gini takes account of the entire distribution, while the Palma / 'top incomes' indices focus on the extreme end(s) of the distribution: a convenient classificatory nomenclature for the two types of indices might be 'across the board' indices and 'tailender' indices respectively.

At a broad level, it seems to be correct to suggest that the 'tailender' approach focuses attention on the normative and policy value of viewing inequality from the standpoint of redistribution—as being most fairly and efficaciously achieved by assessing and bridging the gap between the richest and poorest sections of a population. As a representative member of the class of 'across the board' indices, the Gini coefficient G, by taking stock of not just the upper or lower (or both) tails of a distribution, is uniformly sensitive to the Pigou-Dalton 'principle of transfers', which requires that every progressive rank-preserving pair of income-transfers across the spectrum of a distribution should be reflected in a diminution in the extent of measured inequality.

The concerns of 'tailender' analysts could perhaps be commonly captured in some indicator of the concentration of income with the richest person in a society: a simple example of such an indicator would be the proportionate shortfall of a society's mean income from the income of the richest person, call it R. It is not hard to see that there could be instances of changes in the lower and upper ranges of income in a distribution which leave the Gini coefficient unchanged while affecting the value of R, just as there could be changes in the middle range of incomes which leave R unchanged while affecting the Gini coefficient. For those that would attach specific value to changes in the tails of a distribution as well as to changes in the middle ranges, an exclusive reliance on G would, in the first of the above two instances mentioned, entail losing something of value in failing to have a special focus on the tails; and an exclusive reliance on a measure such as R would, in the second instance, entail losing something of value in neglecting the central parts of a distribution.

In the view outlined above, each of the G and R measures has a distinctively useful property lacking in the other. With Cobham et al (2015; p.17), "…one may conclude that no single measure is likely to meet every concern", and this suggests a case for plurality. But plurality in what sense? In one sense, the requirement would be to present statistics on inequality in terms of *both* the 'across the board' and 'tailender' types of measures—the Gini on the one hand, and on the other, the income-share of the top 10 per cent, the income-share of the bottom 40 per cent, the Palma Ratio, the income-shares of the top 1 per cent and of the top 0.1 per cent, the indicator R, and so on. A second sense, which is the one advocated in the present paper, would to be to go beyond, and to combine features of the two representative types of measures—specifically the G and R measures—in a single composite index, which is here christened D. The derivation of D is very simple: all it requires is an elementary extension of an index of poverty known as the Sen-Shorrocks-Thon measure—see Shorrocks (1995, 1996)—into an index of inequality, along lines which are elaborated on in the following section. The rest of the essay is devoted to an elaboration of the idea just outlined.

An important part of this idea is that there could be specific pairs of income distributions which can be discriminated only in terms of the measure R and not in terms of the Gini coefficient G, and other pairs which can be discriminated only in terms of G and not in terms of R: in such cases, one may wish to proceed according to the ranking effected by a composite measure such as D which accords weight to *both* G and R. It is conceivable that in the event of a tie according to G, D will rank the distributions in the same way as R does; or that in the event of a tie according to R, D will rank the distributions in the same way as G does.

Suppose a *status quo* distribution $\mathbf{x}$ is one that displays an order of inequality which the policy-maker would like to mitigate, and that there are two specific 'improved' distributions $\mathbf{y}$ and $\mathbf{z}$ which can be achieved by redistribution from $\mathbf{x}$, such that the rankings of $\mathbf{y}$ and $\mathbf{z}$ by R and G are mutually opposing: should the policy-maker base their judgement according to 'across the board' or 'tailender' considerations? It is suggested in this paper that the decision should be based on the ranking of $\mathbf{y}$ and $\mathbf{z}$ according to D, which—depending on the specific empirical distributions involved—will sometimes side with G and sometimes with R: what policy move one should adopt to mitigate inequality could thus depend upon how one chooses to measure inequality. Specifically, it could be risky to base one's judgement entirely on the basis of a wholly 'across the board' or wholly 'tailender' approach to the measurement of inequality, as opposed to taking a more comprehensive view of measurement that subsumes both approaches. Even more specifically, this paper suggests that the sorts of global and country-specific distributions one is confronted by are such that the composite measure D would probably lean distinctly on the side of the 'tailender' approach to measurement espoused by the index R.

Against this background, we begin with a derivation of the inequality measure *D*.

2. The *D* Measure of Inequality

2.1 Deriving the Measure

An income distribution is a non-negative, non-decreasingly ordered *n*-vector $\mathbf{x} = (x_1,...,x_i,...,x_n)$ where x_i is the income of the *i*th poorest person in a community of *n* individuals. For every positive integer *n*, $\mathbf{X}_n$ is the set of all *n*-dimensional distributions, and $\mathbf{X} \equiv \bigcup_{n=1}^{\infty} \mathbf{X}_n$ is the set of all conceivable distributions. For every **x**, $n(\mathbf{x})$ is the dimensionality of **x**, and $\mu(\mathbf{x}) \equiv (1/n(\mathbf{x}))\sum_{i=1}^{n(\mathbf{x})} x_i$ is the mean of the distribution **x**. (On occasion, and where there is no ambiguity, the '**x**' in $n(\mathbf{x})$, $\mu(\mathbf{x})$, etc. will be suppressed.) If $\mathcal{R}$ is the set of real numbers, an inequality measure is a mapping $I : \mathbf{X} \to \mathcal{R}$ such that, for every $\mathbf{x} \in \mathbf{X}$, $I(\mathbf{x})$ is a real number signifying the extent of inequality associated with the distribution **x**.

One such measure of inequality is the Gini coefficient, *G*, and there are many ways of writing it. A relatively straightforward way (Sen, 1973) is one which involves the Borda rank-order weighting formula, and is given, for all $\mathbf{x} \in \mathbf{X}$, by:

$$G(\mathbf{x}) = \frac{n(\mathbf{x})+1}{n(\mathbf{x})} - \left(\frac{2}{(n(\mathbf{x}))^2 \mu(\mathbf{x})}\right)\sum_{i=1}^{n(\mathbf{x})}(n(\mathbf{x})+1-i)x_i.$$

To avoid clogging up the notation, the argument **x** in $G(\mathbf{x})$, $n(\mathbf{x})$, $\mu(\mathbf{x})$, etc. will henceforth be suppressed, and we make a beginning with the expression for the Gini coefficient which can be presented in more relaxed notation as:

$$(1)\quad G = \frac{n+1}{n} - \left(\frac{2}{n^2 \mu}\right)\sum_{i=1}^{n}(n+1-i)x_i.$$

For all $\mathbf{x} \in \mathbf{X}$, let $R(\mathbf{x})$ stand for an 'index of income-concentration with the richest person in the distribution', as measured by the proportionate shortfall of the mean income from the richest person's income, so that

$$(2)\quad R = 1 - \mu / x_n.$$

It is straightforward that $G \in [0,(n-1)/n]$ and $R \in [0,(n-1)/n]$.

Consider now a measure of poverty that can be derived from a graph which Shorrocks (1995, 1996) calls the 'poverty-gap profile' or 'deprivation' profile—a graph that was also discovered and discussed by Spencer and Fisher (1992) and Jenkins and Lambert (1997). Without retracing Shorrocks' work in minute detail, it should suffice to note here that his deprivation profile is based on calculations which involve cumulating the income shortfalls of the poor from the *poverty line z*, which is a level of income that serves to separate the poor segment of a population from its non-poor segment. The single deviation we make from the Shorrocks framework is to replace the poverty line *z* by the income x_n of the richest person in the distribution: this serves to postulate a *relative* deprivation profile, entailing calculating the cumulated shortfalls of each person's income from the richest person's income, thereby facilitating a shift in focus from poverty to inequality.

More specifically, given an (ordered) income distribution $\mathbf{x} = (x_1,...,x_i,...,x_n)$, define

$d_i(\mathbf{x};i/n) \equiv (x_n - x_i)/x_n,\ \ i = 1,...n;$ and

$$D_i(\mathbf{x};i/n) \equiv (1/n)\sum_{j=1}^{i} d_j(\mathbf{x};j/n),\ \ i = 1,...n.$$

The *relative deprivation*, or *inequality*, profile is the graph of $D_i(\mathbf{x};i/n)$, and is obtained by plotting, within the unit square, the set of points $\{(0,0),(1/n,D_1),(2/n,D_2),...,((n-1)/n,D_{n-1}),(1,D_n)\}$, and joining the points by straight lines to obtain a piece-wise linear curve. The profile is an increasing, concave curve which flattens out and becomes horizontal at the penultimate point $((n-1)/n,D_{n-1})$ of the curve (note that $D_n = D_{n-1}$, since $d_n = (x_n - x_n)/x_n = 0$). (More precisely, the curve will flatten out at the point corresponding to the ordinate *q/n* if the richest *n-q* individuals shared the same income.) The 'line of maximum inequality' would represent a situation in which $x_i = 0 \forall i = 1,...,n-1$ and $x_n > 0$, so that $D_i = i/n \forall i = 1,...,n-1$ and $D_n = D_{n-1} = (n-1)/n.$ When *n* is large, the line of maximum inequality can be approximated by the diagonal of the unit square, and the inequality profile by a smooth, concave curve. Figure 1 presents the inequality profile for a situation in which *n* assumes a finite value of 5, and Figure 2 the inequality profile for a situation in which *n* is 'large'. Finally, it should be clear that $D_n = 1 - \mu / x_n \equiv R.$

A well-defined normalized measure of inequality—call it D—can be obtained as a ratio of the area under the inequality profile (call it A) to the area under the line of maximum inequality (call it $A_{\max}$). For a piece-wise linear inequality profile, the area under the curve can be computed by the 'trapezoidal method', which would entail calculating and summing the areas of a number of triangles and rectangles. Some tedious algebra, which we shall here avoid by relegating it to an Appendix, should confirm that an expression for the measure we are after can be written as follows. For all $\mathbf{x} \in \mathbf{X}$ (and after suppressing the argument **x**):

(3a) $D = [n^2/(n^2-1)][1-(\mu/x_n)(1-G)]$ for finite n; and

(3b) $D \to 1-(\mu/x_n)(1-G)$ as $n \to \infty$.

or, making the appropriate substitution of R for $1-\mu/x_n$ from Equation (2),

(4a) $D = [n^2/(n^2-1)][R+(1-R)G]$ for finite n; and

(4b) $D \to R+(1-R)G$ as $n \to \infty$.

Notice from (4b) that the inequality index D is a function of both the 'across the board' measure G and the 'tailender' measure R, and is non-decreasing in each of its arguments: $D = D(R,G)$; $\partial D/\partial R = 1-G \geq 0$, and $\partial D/\partial G = 1-R \geq 0$.

A final comment may be in order. This has to do with the link between an inequality index and a poverty index, and how the former may be derived from the latter. On the connection between the Gini inequality coefficient and the Sen poverty index, Sen (1976; pp.225-226) says: "The problem of measurement of inequality and that of poverty can be seen to be two intertwined exercises. The measure of inequality corresponding to the [Sen] measure of poverty P can be defined in the following way: …The measure of inequality η corresponding to the poverty index P … is the value obtained in place of P by replacing q (the number of poor) by n (the total number of people in the community), and replacing z (the poverty level) by [μ] (the mean income of the community)….*The measure of inequality η corresponding to the poverty index P is the Gini coefficient* [G] *for large n*" [emphasis in original]. It is literally undeniable, simply as an arithmetical truth, that when z is replaced by μ and q by n, then the Sen poverty index converges on the Gini coefficient for large n—but this begs the question of how one may confer a 'physical' interpretation on the statement: when the poor are defined as those with incomes less than the poverty line (the weak definition of the poor), then a situation in which everybody is poor in relation to a poverty line which is the average income implies that every person must have a below-mean income, an impossibility; and when the poor are defined as those with incomes not exceeding the poverty line (the strong definition), the situation in question is one in which everybody has the mean income, an outcome which coincides with the special case of a zero-value for the Gini coefficient. These problems do not arise if we start out with the Sen-Shorrocks-Thon poverty index, and derive an inequality measure by replacing the poverty line z by the richest person's income x_n (in which case, by the strong definition of poverty, since no-one can have an income exceeding the richest person's income, it must be the case that q will be replaced by n): for large n, the resulting inequality index, as we have seen, will be given by:

$D = 1 - (\mu/x_n)(1\text{-}G)$.

2.2 Some Properties of D

Here we take brief stock of some properties of the inequality measure D (the discussion is brief, because these properties, discussed in earlier literature on the Sen-Shorrocks-Thon poverty index—see Shorrocks, 1995, 1996—carry over naturally from the poverty index to the inequality index). Some commonly proposed axioms for inequality measurement are the following. *Symmetry* requires the value of the inequality measure to be invariant to permutations of incomes across individuals ('personal identities should not matter'); *continuity* requires the measure to be continuous in incomes ('similar income distributions should display similar inequality values'); *normalization* requires the inequality measure to assume a value of zero when incomes are perfectly equally distributed; the *transfer axiom* requires measured inequality to increase in the face of a rank-preserving regressive transfer, that is, a transfer from one person to another who is richer; *scale-invariance* requires the measure to remain unchanged when all incomes in a distribution are increased or decreased in the same proportion ('mean independence', or 'neutrality to units of measurement'); and replication-invariance requires inequality to remain unchanged in the presence of a population replication, that is, when the number of persons at each income level is increased by an identical factor of k, where k is any positive integer ('inequality depends only on relative, not absolute, population frequency'). As it happens, and as can be easily verified, the inequality measure D satisfies all of the preceding standard axioms of inequality measurement.

A further property one may mention is that of *weak transfer-sensitivity*, which requires that a given progressive transfer of income between two individuals which is accompanied by an identical regressive transfer between a pair of richer individuals should cause the inequality measure to decline in value whenever the individuals in each pair of persons involved in the transfers are separated by the same income and the same number of income-ranks. It is well known that the Gini coefficient G violates this requirement: it will remain unchanged by the transfers just described. On the other hand, if one of the individuals in the richer pair of persons involved in the transfer should happen to be the richest individual in the distribution, then while G will remain unchanged, the measure R will register a rise in value: consequently, in this special case, the index $D \equiv R + (1 - R)G$ will also increase. This would amount to an instance of '*reverse* transfer-sensitivity', but it is worth asking if this is a necessarily perverse outcome: is one's moral intuition unambiguously clear that it is better to tax one person to benefit a poorer one than to tax a richer person to benefit the richest? We shall leave it at that, as a query on the unqualified appeal of the weak transfer-sensitivity property.

2.3 The Measure D as a Tie-Breaker

A consequence (suggested in the introductory section) of combining G and R in a composite index such as D is that, in specific cases, when inequality in any two distributions cannot be distinguished by any one of the two measures G and R, the other measure can play a decisive role in breaking the tie registered by the first. A simple numerical example should help. Consider the following three 5-person distributions, each of which has a mean of 30:

$\mathbf{x} = (10,20,30,40,50)$;

$\mathbf{y} = (10,20,25,45,50)$; and

$\mathbf{z} = (15,15,25,40,55)$.

Distribution **y** can be seen to have been derived from **x** by a regressive transfer of income from person 3 to person 4 in the middle group; and distribution **z** can be seen to have been derived from **y** through a progressive transfer of 5 units of income from person 2 to person 1, accompanied by a regressive transfer of 5 units from person 4 to the richest person 5. Table 1 presents computed values, for each of the distributions, of the inequality indices *G*, *R* and *D* (calculated in accordance with the expressions for these indices furnished in Equations (1), (2) and (4a) respectively).

Table 1: Inequality Values for the Distributions x, y and z

Distribution↓/ Inequality Measure→	*G*	*R*	*D*
x	0.2667	0.4000	0.5833
y	0.2800	0.4000	0.5917
z	0.2800	0.4545	0.6325

Source: Authors' calculations

Concerned as it is only with the mean in relation to the richest person's income (both of which remain unaffected by the transition from **x** to **y**), the *R* measure is unable to distinguish between **x** and **y**, whereas the Gini, by taking stock of the within-(middle-)group redistribution of incomes, pronounces that **y** has more inequality than **x**, a judgement that is supported by the *D* measure.

Similarly, in moving from **y** to **z**, we have one progressive and one regressive transfer, and the progressive transfer is seen as being neutralized by the regressive transfer by the Gini coefficient which pronounces **y** and **z** as having the same extent of inequality, whereas the measure *R*, with its pronounced emphasis on what happens to the mean in relation to the richest person's income, registers the fact that the gap between the two has risen in **z** vis-à-vis **y**, and in this judgement, it is supported by the ranking of **y** and **z** by the composite *D* measure.

Thus, in specific cases, the composite index *D* serves the purpose of a tie-breaker, and assists in taking a fuller account of inequality than would be the case if one relied only on one or the other of its component measures *G* and *R*. That is to say, *D* serves the purpose of a consistent real-valued measure of inequality that takes stock of the useful properties which each of two candidate measures possesses in such a way that, in temporal inequality comparisons for example, the trend displayed by *D* could differ from that displayed by either of the candidate measures. (In the preceding simple numerical illustration, if **x**, **y** and **z** are interpreted as the distributions valid for three successive years 1, 2 and 3 respectively, then it can be seen that Gini increases from year 1 to year 2 and remains the same from year 2 to year 3, and R remains the same from year 1 to year 2 and rises from year 2 to year 3, whereas D increases continuously from year 1 to year 2 to year 3.)

The preceding observations are ones which hold 'in principle', and can be illustrated by simple numerical examples, such as the ones we have employed. Of obviously greater interest would be the patterns of uniformity or otherwise displayed by each of G, R and D in the actual sets of time-series and cross-section empirical income or wealth distributions which one is confronted by in the real world. What are the differences among G, R and D in the range of variations over time or across space that these measures exhibit, and how are these differences to be interpreted? In specific policy choices for the mitigation of inequality, is there a case for being guided by D, and if so, is there a case for seeing one or other of G and R as being more 'decisive' in 'influencing' D? Some of these issues are addressed in what follows, through illustrative empirical applications to inequality in the world and in India.

3. Some Empirical Illustrations of Aspects of Income Inequality

3.1 Elements of the Global and Regional Picture

The website of the World Inequality Database (WID) is a repository of information on several indicators of income and wealth inequality, globally and across regions and countries of the world. Among other things, the WID has a regionally disaggregated time-series on the per adult (pre-tax) income-share of the richest 1 per cent of the population. For any given year and geographical entity, suppose μ to stand for mean income, $\mu_{0.01}$ for the mean income of the richest 1 per cent, and $s_{0.01}$ for the income share of the richest 1 per cent. Then, by definition, $0.01\mu_{0.01}/\mu = s_{0.01}$, whence $\mu/\mu_{0.01} = 0.01/s_{0.01}$: if we treat the average income of the richest 1 per cent of the population as a proxy for the richest person's income, then an estimate of the measure R (see Equation 2) is provided by the quantity $1-(0.01/s_{0.01})$. Needless to say, this is a very gross assumption to make since, typically, one must expect quite substantial inequality in the distribution of incomes within the top 1 per cent of the population, so the quantity $1-(0.01/s_{0.01})$ must serve as a considerable under-estimate of the actual value of *R*. But we will make do with this in the absence of information on the richest person's income. The WID also provides information on the Gini coefficient *G* of inequality in the pre-tax distribution of adult income. Table 2 presents estimates of the income-share of the poorest 1 per cent ($s_{0.01}$), our derived value of the measure *R* ($1-(0.01/s_{0.01})$), and the Gini coefficient *G*, based on information in the WID, for two points in time in our recent history, 1980 and 2020.

Table 2: Some Global and Regional Indicators of Inequality: 1980, 2020

Region	$S_{0.01}$ (1980)	R (1980)	G (1980)	D (1980)	$S_{0.01}$ (2020)	R (2020)	G (2020)	D (2020)
World	0.1679	0.9404	0.69	0.9815	0.1978	0.9524	0.67	0.9843
East Asia	0.1983	0.9496	0.66	0.9829	0.1760	0.9375	0.59	0.9744
Europe	0.0819	0.8790	0.43	0.9310	0.1217	0.9167	0.48	0.9567
Latin America	0.1875	0.9467	0.70	0.9840	0.2240	0.9615	0.72	0.9892
Middle East & North Africa	0.3372	0.9703	0.75	0.9926	0.2302	0.9565	0.66	0.9852
North America	0.1025	0.9024	0.45	0.9463	0.1804	0.9474	0.57	0.9774
Central Asia	0.1356	0.9263	0.57	0.9683	0.1866	0.9500	0.63	0.9815
South & Southeast Asia	0.1792	0.9442	0.54	0.9743	0.2064	0.9500	0.61	0.9805
Sub-Saharan Africa	0.1970	0.9492	0.70	0.9848	0.1953	0.9474	0.66	0.9821

Note: $s_{0.01}$ stands for the income-share of the top 1 per cent; *R* stands for the quantity $1-\mu/\mu_{0.01}=1-0.01/s_{0.01}$; *G* stands for the Gini inequality coefficient; and *D* stands for the inequality measure *R* + (1-*R*)*G*.

Source: Figures based on information from the World Inequality Data Base (Home - WID - World Inequality Database: https://wid.world).

Among other things, Table 2 reveals that the Gini coefficient displays a wide range of variation, both across regions at a given point of time, and for given regions across time. Thus, *G* varies from 0.43 for Europe to 0.75 for MENA in 1980, and from 0.48 for Europe to 0.72 for North America in 2020; and one can see a significant reduction in *G* for MENA between 1980 and 2020 (from 0.75 to 0.66), together with significant increases for North America (from 0.45 to 0.57), for Central Asia (from 0.57 to 0.63), and for South and Southeast Asia (from 0.54 to 0.61). By contrast, there is little variation in the values of the measures *R* and *D*, whether we take a 'cross-section' or 'time-series' perspective: in fact, with the exception of Europe in 1980, all *R*-values for all regions are at least 0.90 in both 1980 and 1990; and the *D*-values for most regions are as close to the maximal value of unity as makes no difference! As to what to make of this, we shall return to the issue after noting a similar pattern for India.

2.4 Indian Income Inequality

Bharti, Chancel, Piketty and Somanchi (2024) is an important historical profile of income and wealth inequality in India. Estimates of top incomes for India have been provided in the study at the very fine-grained level of the richest 0.1 per cent of the population. As before, letting μ stand for average income, $\mu_{0.001}$ for the average income of the top 0.1 per cent, and $s_{0.001}$ for the income-share of the top 0.1 per cent, we have: $\mu / \mu_{0.001} = 0.001 / s_{0.001}$; and if we treat the average income of the richest 0.1 per cent as a proxy for the richest person's income—this would be a better approximation than the average income of the richest 1 per cent, though still an under-estimate—then an estimate of the measure *R* would be given by the quantity $1 - 0.001 / s_{0.001}$. Using the information on the income-shares of the top 0.1 per cent available in Bharti et al (2024), and WID information on the Gini coefficient for the distribution of adult income in India, we can generate estimates of the inequality indices *G*, *R* and *D* for India, at decadal intervals, over the 70-year period 1952 to 2022. The relevant statistics are provided in Table 3.

As in the case of Table 2, Table 3 also suggests that the Gini is a discriminating measure of inequality: the figures in the Table indicate a noticeable decline in *G* from 1952 to 1982, and an even more perceptible increase from 1982 to 2022. In contrast, the indices *R* and *D*—though they preserve the over-time ranking according to *G*—display little in the way of year-to-year variation: *R* is compressed in the narrow range of 0.9412 (in 1982) to 0.9896 (in 2022), and *D* in the even narrower range of 0.9647 (in 1982) to 0.9959 (in 2022); furthermore, these are ranges whose lower limits themselves come close to describing complete inequality. This seems to call for some analysis and comment, which are attempted in what follows.

Table 3: Some Indicators of Income Inequality for India: 1952 – 2022

Year	$s_{0.001}$	R = 1 - 0.001/$s_{0.001}$	G	D = R + (1-R)G
1952	0.047	0.9787	0.46	0.9885
1962	0.046	0.9783	0.45	0.9881
1972	0.036	0.9722	0.42	0.9868
1982	0.017	0.9412	0.40	0.9647
1992	0.028	0.9643	0.44	0.9800
2002	0.054	0.9815	0.49	0.9906
2012	0.082	0.9878	0.60	0.9951
2022	0.096	0.9896	0.61	0.9959

Note: $s_{0.001}$ stands for the income-share of the top 0.1 per cent; R stands for the quantity $1-\mu/\mu_{0.001}=1-0.001/s_{0.001}$; G stands for the Gini inequality coefficient; and D stands for the inequality measure $R + (1-R)G$.

Source: Figures based on information from the Data Appendix: 'Table B1: Per-adult pre-tax national income shares (%), 1951-2022' in Bharti et al (2024), and the World Inequality Data Base (Home - WID - World Inequality Database: https://wid.world).

4. Attempt at an Interpretation

The patterns thrown up by Tables 2 and 3 provoke one to ask if the measures R and D are essentially blunt instruments which fail to discriminate amongst income distributions by displaying a tendency to judge extreme inequality in all distributions. Is there something 'wrong' with these measures? It is true that R fails to satisfy even the elementary Pigou-Dalton transfer axiom, and this has been held against 'tailender' inequality indices such as the Palma Ratio. However, this is not a charge that can be levelled against the measure D which incorporates both R and G within itself: indeed, as we have seen in Section 2.2, D seems to tick all the boxes in terms of the basic, approved axioms of inequality measurement: symmetry, continuity, normalization, transfer, and scale- and replication-invariance.

Perhaps, then, the 'fault' is not with the measure D but with the sorts of distributions whose inequality D is called upon to assess! There cannot, after all, be much that is specially 'right' with a distribution in which the average income of a society is barely 5 per cent of the average income of a miniscule 0.1 per cent of its richest members—and these are the sorts of distributions featured in Tables 2 and 3: surely, these distributions did not *have* to be like this! It also bears remarking that there is perhaps not much 'surprise value' occasioned by such inequality: it is something to which we may have just got accustomed, something that has got 'normalized' over time as an essential feature of the economic systems by which we are governed. Indeed, and regarding the matter from a semantic point of view, the notion of what is 'normal' is frequently conflated with the notion of what is 'common', and there is a case for guarding against such conflation, however natural the impulse to yield to it may be. For instance, measured iron deficiency in a human *is* a pathology, notwithstanding the possibly wide prevalence of anaemia in a society.

That what needs to be fixed is not so much a particular measure of inequality as the distribution whose inequality is being assessed can be illustrated with a counterfactual example. The object of the exercise is to postulate an initial ('status quo') distribution **a**, based on the actual 2020 global distribution of income; to consider two less inequitable distributions, call these **b** and **c** respectively; to ask which of these distributions is intuitively more appealing; and to examine which sort of inequality measure—an 'across the board' measure such as G or a 'tailender' measure such as R—would support one's intuition.

Against this background, consider some summary inequality indicators for the global distribution of income in 2020, which are available in Chancel and Piketty (2021). We note the following: from Table 5 of the cited paper, the income-share of the top 10% is 55%; from Table 6, the income-share of the bottom 50% is 7%; we can therefore deduce the income share of the middle 40% to be 38% [100% – (7%+55%)]; from Table 8, the income-share of the top 1% is 21%[2], whence the income-share of the 90%-99% segment of the population must be 34% [the income-share of the top 10%, namely 55%, less the income-share of the top 1%, namely 21%]; Figure 8 of the cited paper indicates that the average income of the global top 0.1% is 602 times higher than the average income of the bottom 50%, and since we know the income share of the bottom 50% to be 7% from Table 6, we can deduce the income-share of the top 0.1% to be (0.001)(0.14)(602), or 8.43%; and deducting this figure (8.43%) from the income-

[2] The WID data employed in Table 2 suggest a slightly lower figure of 19.78%. Our concern, however, is to provide a simple illustrative example, and not with fineness of detail.

share of the top 1% (21%) yields the income-share of the 99%-99.9% segment of the population, namely 12.57%. We can employ these data to generate some ordinates of the Lorenz curve, in terms of the cumulated income-shares corresponding to a set of cumulated population-shares, as in Table 4.

Table 4: Some Ordinates of the Lorenz Curve for the Global Distribution of Income in 2020

Cumulated Population Shares	Cumulated Income Shares
Poorest 0%	0%
Poorest 50%	7%
Poorest 90%	45%
Poorest 99%	79%
Poorest 99.9%	91.57%
Poorest 100%	100%

Source: Authors' calculations based on data in Chancel and Piketty (2021).

These are awkward, irregular size-class intervals, involving only a few observations, but they do provide some information for calculating the Gini coefficient: the trapezoidal approximation formula (which would surely be an underestimate of the true value) yields a Gini of 0.6281 for the 2020 global distribution of income (the WID data suggest a higher value of 0.67). Taking the average income of the richest 0.1% to be a proxy for the richest person's income yields a value for the measure R of 0.9881, and given that G is 0.6281, we obtain a value for D ($=R + (1-R)G$) of 0.9956.

The inequality described in the preceding numbers can be equivalently expressed as a particular distribution of a thousand dollars among a thousand individuals, where it is assumed that within each size-class of income percentiles, income is equally divided. This distribution—call it the *status quo* distribution **a**—is as follows:

Distribution **a**:
Income of each of the 500 poorest persons: $0.14;
Income of each of the next 400 poorest persons: $0.95;
Income of each of the next 90 poorest persons: $3.78
Income of each of the next 9 poorest persons: $13.97
Income of the richest person: $84.30

As we have seen, $G(\mathbf{a}) = 0.6281$; $R(\mathbf{a}) = 0.9881$; and $D(\mathbf{a}) = 0.9956$.

Suppose one had the option of arranging a less inequitable distribution of incomes, and that exactly two alternative distributions—**b** and **c**—are available to the arranger. The distributions, with accompanying values of the inequality indices G, R and D, are described below:

Distribution **b**[3]:
Income of each of the 999 poorest persons: $0.9166
Income of the richest person: $84.30
$G(\mathbf{b}) = 0.0833;\ R(\mathbf{b}) = 0.9881;$ and $D(\mathbf{b}) = 0.9891.$

Distribution **c**:
Income of each of the 500 poorest persons: $0.85;
Income of each of the next 400 poorest persons: $1.02;
Income of each of the next 90 poorest persons: $1.50
Income of each of the next 9 poorest persons: $3.00
Income of the richest person: $5.00
$G(\mathbf{c}) = 0.1025;\ R(\mathbf{c}) = 0.8000;$ and $D(\mathbf{c}) = 0.8205$.

In moving from the status quo distribution **a** to the distribution **b**, the income-share of the top 0.1% has been retained at 8.43% (so that each person in the richest 0.1% receives an income of $84.30), while the remaining income is divided equally among the remaining 999 persons, so that each of them receives a paltry income of just under 92 cents. In the transition from **a** to **c**, the income-share of the top 0.1% is drastically reduced from 8.43% to 0.5% (nevertheless allowing the average income of the top 0.1% to be five times the mean income), and while the top 10% fare worse in **c** than in **a**, it is the other way around with the bottom 90%. Between **b** and **c**, the bottom 50% fare a little worse in **c** than in **b**, the next 49.9 % fare better, and of course the top

[3] It is useful here to note a general result: in a situation (as described by distribution **b**) where the poorest n-1 individuals share the same income, the Gini coefficient is given by the difference between the income-share of the richest individual (s_n) and his share in the population ($1/n$). In the instant case, $s_n = 0.0843$ and $1/n = 0.001$, whence $G = s_n - 1/n = 0.0833$, as recorded. (Readers interested in a precise derivation of the result just stated are referred to Alvaredo (2011): this result can be obtained as a special case from Alvaredo's Equation (4)—which is a general expression for the Gini coefficient G—by setting the quantities G^* and G^{**} in his Equation equal to zero. It may be added that Alvaredo's concern in the cited paper is to show that the failure of surveys to capture the under-reporting of top incomes can cause the magnitudes and trends of the Gini coefficient to be considerably more flattering than they actually are, even if the population share of top income holders should be infinitesimal, an outcome that would hold *a fortiori* when this population share is finite.)

0.1% are cut down substantially to size. It does appear, taking everything into account, that **c** avoids the really rather abject picture presented by **b** of 999 persons out of a thousand being cut off with a very low income in order to accommodate the richest person's disproportionately huge income. Yet, as it happens, the 'across the board' inequality measure *G* favours **b** over **c**, in contrast to the ranking by the 'tailender' measure *R*, whose judgement is also endorsed by the composite measure *D*.

The judgement in question is resistance to extreme concentration at the top of the income distribution. Figure 3 plots the behaviour of the measure *R* in response to an increase in the income-share $s_{0.001}$ of the top 0.1%: as can be seen, *R* rises very sharply with initial increases in $s_{0.001}$, so that, by the time the income-share of the richest 0.1% has reached 1%, *R* has already attained a value of 0.9, not far from its maximum value of unity; thereafter, *R* rises very slowly and asymptotically toward its ceiling-value. Since *R* for any distribution defines the lower bound on the composite measure *D*, the latter also ceases to respond very sensitively to increases, beyond a point, in the income-share of the richest. This echoes an old Tamil proverb: 'If you are drowning, how does it matter if the water level is a foot above your head or a yard?'[4]

To summarize: *D* seems to be a plausible measure of inequality, one which, given the sorts of income distributions that obtain in the world today, can be interpreted as endorsing the capping of top incomes as a means of mitigating inequality. The issue is addressed in the following section.

5. Curbing Top-Heaviness

The *R*-value of the global income distribution in 2020, at 0.9881, can be brought down to 0.90 (still a high value!), by having an average income for the richest 0.1% which is 10 times the global average income. This is no mean factor of difference, except when it is compared to the prevailing standard, by which the average income of the top 0.1% is in excess of 84 times the average global income. There is a case for not

[4] Thanks to Professor Venkatesh Athreya for the translation.

yielding too easily, from simple accustomedness to such grotesque prevailing standards.

What is being suggested is the necessity of capping top incomes. This would conform with a distributional ethic characterized by Robeyns (2019) as 'limitarianism'. As she puts it (Robeyns, 2019; pp. 252-253): "Limitarianism as an ethical or political view is, in a certain sense, symmetrical to the view that there is a poverty line and that no one should fall below this line. Limitarianism claims that one can theoretically construct a riches line and that a world in which no one would be above the riches line would be a better world."

Robeyns proposes two arguments in support of limitarianism, the 'democratic argument' and the 'urgent unmet needs argument'. The former argument can be summarized in her own words (Robeyns, 2021; p.254): "The first justification for the limitarian view relates to democracy and the worry that massive inequalities in income and wealth undermine the value of democracy and the ideal of political equality in particular." The second argument is predicated on the requirement of financial resources, presumably to be supplied by the super-rich, in a world in which there are urgent unmet needs arising from at least one of three conditions: "extreme global poverty", "local or global disadvantages", and severe "collective action problems", all of which are held by Robeyns, and reasonably so, in our view, to obtain in the world as we know it (Robeyns, 2019).

One approach to capping the incomes of the very rich would be guided by the perceived need of eradicating poverty. This would be compatible even with the conservative distributional doctrine of '*sufficientarianism*' due to Frankfurt. As he puts it (Frankfurt, 1987; p.21): "With respect to the distribution of economic assets, what *is* important from the point of view of morality is not that everyone should have *the same*, but that each should have *enough*. If everyone had enough, it would be of no moral consequence whether some had more than others. I shall refer to this alternative to egalitarianism—namely, that what is morally important with respect to money is for everyone to have enough—as 'the doctrine of sufficiency'." Hassoun (2021) points out that any theory of rights and justice, and indeed of common decency and human dignity, must insist on securing a 'minimally good life' for all members of a society. She notes that sufficientarianism is an unlikely candidate for an adequate theory of equality, but that it does address a necessary part of such a theory, in the sense and to the extent that it does not tolerate inequality in the presence of poverty interpreted as a condition of ‘not having enough’.

Seen in this light, sufficientarianism may be seen as endorsing a limited form of limitarianism in the following way. Suppose z to be some conventional ‘poverty line’ level of income and that there are q poor individuals out of a total of n persons in the society under review. Let Q stand for the ‘aggregate poverty deficit’, that is, the total income that would be needed to raise all the poor persons to the poverty line z. One way of meeting the poverty deficit would consist in a system of taxation of the richest of the rich wherein one first reduces the richest person’s income to the level of the next richest person’s income; if the income thus taxed covers the poverty gap, the exercise can be terminated; if not, the incomes of the two richest persons are reduced to the income of the third richest person;…; and so on, until one reaches that marginal person, with an income, say, of x^*, with whom the total tax raised from, say, the richest r individuals, just bridges the required poverty shortfall. The post-tax-cum-transfer distribution of income will be one in which the poorest q persons share an income of z each, the richest r individuals share an income of x^* each, and the middle $n-(q+r)$ persons have the same incomes as in the pre-tax-cum-transfer distribution. It can be shown that the resulting distribution cannot be Lorenz-dominated by any other system of redistributive taxation aimed at ensuring no more than that every poor person in the status quo distribution is enabled to just escape poverty. This system of taxation and transfers is one which has been advanced by Jayaraj and Subramanian (1996, 2010), and related approaches can be found in Medeiros (2006), and Basu (2024). The point to note is that the prescription is compatible with the demands of sufficientarianism and entails capping the richest person’s income at the level x^*, derived along the lines just described.

The level at which richest incomes are capped need not be dictated by considerations only of poverty eradication. The upper limit on individual richness could be prescribed in terms of some independent normative judgement on how far the inequality measure R may be permitted to rise. So, for example, let $\hat{x}$ be that level of income shared by the richest 0.1% such that the income-share of this segment of the population is 1% (or, equivalently, also the level of income at which the value of the inequality measure *R* is 0.90). (One may of course wish to choose an even lower upper limit for *R*, in which case $\hat{x}$ would also be lower.) A version of limitarianism which takes stock of sufficientarianism only to the extent that poverty-eradication is seen as an essential but not necessarily adequate aspect of distributional fairness may require that the richest person's income should be capped at a level $\tilde{x}$ given by $\tilde{x} \equiv \min(\hat{x}, x^*)$. The tax collected from capping the incomes of the richest 0.1% at $\tilde{x}$ can be used to bridge the aggregate poverty deficit, or to contribute to a universal basic income, or to enhance budgetary provisions for social sector spending on, say, education or health. What is integral to a desired distribution is the capping of top incomes.

Objections to capping may be based on the fear that redistributive taxation may result only in generalized immiserization of a society. This seems to be a misplaced fear, even in the context of countries with a fair amount of poverty in them, such as India. Income streams are strongly correlated with wealth stocks, and the massive asset holdings of a miniscule proportion of the ultra-rich in India point to the enormous potential for substantial redistribution through mild wealth taxation of an order that should not occasion any serious reversal of fortune for the super-wealthy (Subramanian, 2024). There is then the argument that taxes are distortionary, inefficient, and ultimately injurious to the incentives that drive the very rich. But this argument does not seem to take into account that incentives work for the poor as well as for the rich: it is not being advocated that the proceeds from taxation be dumped in the ocean, but rather that they be redistributed to the poor so that, surely, negative incentives from taxes at the upper tail of the distribution may be compensated by positive incentives from transfers at the lower tail. A third objection that is frequently resorted to relates to how taxes on the very wealthy will be avoided or evaded or responded to with the flight of capital to more tax-friendly destinations. The present paper is not oriented to an examination of the fine-tuned details of resistance-proof mechanisms of taxation, but the literature is not lacking in scrupulous investigation of these and related issues: three examples would be Cobham, Faccio, Garcia-Bernardo, Jansky, Kadet and Picciotto (2022), Zucman (2024) and Saez and Zucman (2024). The enormous scepticism underlying the view that taxation of top incomes is inevitably destined to fail is pushed back in the instructive view of Saez and Zucman (2024; p. i1130) that: "Tax avoidance, tax competition and tax evasion are not laws of nature."

In any event, it is not clear how the untrammelled growth of top incomes beyond their already cosmic contemporary levels can continue to be seen as a viable proposition (on the urgency of taking corrective action against global economic inequity, see Ghosh, Ocampo and Stiglitz, 2023). We are not unaware that this position can be, and for the most part is, viewed as a counsel of perfection; it is just that we believe it would be more helpful, and also realistic, to view it as a counsel against serious perversion.

6. Concluding Observations

'Across the board' measures of inequality such as the Gini coefficient *G* take account of the entire distribution of incomes, while 'tailender' measures such as the *R* index focus attention on income-concentration at the tail(s) of the distribution. The debate on these alternative approaches to inequality measurement has tended to adopt an 'either/or' position in the matter. In this paper, by contrast, the emphasis has been on a form of plurality which affords place for both an 'across the board' and a 'tailender' measure—specifically for both *G* and *R*—by combining them in a composite indicator of inequality *D* which might have something to commend it as a compromise measure, one which affords 'voice' to each of its component indices without stifling the other. There is little that is original in the derivation of the index *D*: it falls out naturally from a mild extension of the Sen-Shorrocks-Thon measure of poverty.

In assessing inequality in distributions with high 'top income' concentrations, such as have obtained for decades, both globally and in individual regions and countries, it appears that the measure *D* is decisively influenced by the 'tailender' component *R*. Reducing *D* would be fundamentally dependant on reducing *R* to reasonable levels. This points in the general direction of the absolute importance of capping top incomes in distributions with out-of-control top income shares. Both morality and measurement suggest that there is good reason to support the distributional doctrine of limitarianism.

References

Alvaredo, F. (2011). 'A Note on the Relationship Between Top Income Shares and the Gini Coefficient,' *Economics Letters*, 110(3): 274-277.

Alvaredo, F., A. B. Atkinson, T. Piketty and E. Saez (2013). 'The Top 1 Percent in International and Historical Perspective,' *Journal of Economic Perspectives*, 27(3): 3-20.

Atkinson, A. B., T. Piketty and E. Saez (2011). 'Top Incomes in the Long Run of History,' Journal of Economic *Literature*, Anthony B. Atkinson, Thomas Piketty, and Emmanuel Saez, 49(1): 3-71.

Basu, K. (2024). *Reason to be Happy: Why Logical Thinking Can Lead to a Better Life*. Penguin Books: UK.

Bharti, N. K., L. Chancel. T. Piketty and A. Somanchi (2024). 'Income and Wealth Inequality in India, 1922-2023: The Rise of the Billionaire Raj,' WP No. 2024/08, World Inequality Lab.

Chancel, L. and T. Piketty (2021). 'Global Income Inequality, 1820-2020: The Persistence and Mutation of Extreme Inequality,' *Journal of the European Economic Association*, 196(6): 3025-3062.

Cobham, A., Schlogl, L. and Sumner, A. (2015). "Inequality and the Tails: The Palma Proposition and Ratio Revisited," DESA Working Paper No. 143 ST/ESA/2015/DWP/143.

Cobham, A., T. Faccio, J. Garcia-Bernardo, P. Jansky, J. Kadet and S. Picciotto (2022). 'A Practical Proposal to end Corporate Tax Abuse: METR, A Minimum Effective Tax Rate for Multinationals,' *Global Policy*, 13(1): 18-33.

Ghosh, J., J. A. Ocampo and J. E. Stiglitz (2023). 'Put the UN in Charge of International Taxation,' *Project Syndicate*, November 23, 2023. Available at: Put the UN in Charge of International Taxation by Jayati Ghosh, José Antonio Ocampo and Joseph E. Stiglitz - Project Syndicate

Frankfurt, H. (1987) 'Equality as a Moral Ideal,' *Ethics*, 98(1): 21-43.

Hassoun, N. (2022). 'Good Enough? The Minimally Good Life Account of the Basic Minimum,' *Australasian Journal of Philosophy*,100(2): 330-341.

Jayaraj, D. and S. Subramanian (1996). `Poverty Eradication Through Redistributive Taxation: Some Elementary Considerations', *Review of Development and Change*, 1(1): 73-84, revised version in D. Jayaraj and S. Subramanian (2010): *Poverty, Inequality, and Population: Essays in Development and Applied Measurement*. Oxford University Press: Delhi.

Jenkins, S. P. and P. J. Lambert (1997). 'Three I's of Poverty Curves, with an Analysis of UK Poverty Trends,' Oxford Economic Papers, 49(3): 317-327.

Milanovic, B. (2015). "Why I am still not excited by the Palma Index," available at: https://glineq.blogspot.com/2015/06/why-i-am-still-not-excited-by-palma.html

Medeiros, M. (2006). 'The Rich and the Poor: The Construction of an Affluence Line from the Poverty Line', *Social Indicators Research*, 78(1): 1–18.

Palma, G. (2011). "Homogeneous Middles vs. Heterogeneous Tails, and the End of the 'Inverted-U': The Share of the Rich is What It's All About," *Development and Change*, 42(1), 87–153.

Robeyns, I. (2019). 'What, if Anything, is Wrong with Extreme Wealth?,' *Journal of Human Development and Capabilities*, 20(3): 251-266.

Saez, E. and G. Zucman (2024). 'Top Incomes and Tax Policy,' *Oxford Open Economics*, 3: i1130-i1132.

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

Sen, A. (1976). 'Poverty: An Ordinal Approach to Measurement,' *Econometrica*, 44(2): 219-31.

Shorrocks, A. F. (1995). 'Revisiting the Sen poverty index,' *Econometrica, 63*(5): 1225–1230.

Shorrocks, A. F. (1996). 'Deprivation profiles and deprivation indices,' In S. Jenkins, A. Kapteyn, & B. Vaan Praag (Eds.), *The Distribution of welfare and household production: International perspectives*. London: Cambridge University Press.

Spencer, B. D. and S. Fisher (1992). 'On Comparing Distributions of Poverty Gaps,' *Sankhya, Series B*, 54(1): 114-126.

Subramanian, S. (2024). 'The Spirit Underlying Some Recent Proposals for Redistributive Taxation,' The India Forum, July 2. Available at: The Spirit Underlying Some Recent Proposals for Redistributive Taxation | The India Forum

Zucman, G. (2019). 'Global Wealth Inequality,' *Annual Review of Economics*, 11: 109-138.

Figure 1: The inequality profile for finite n (n = 5)

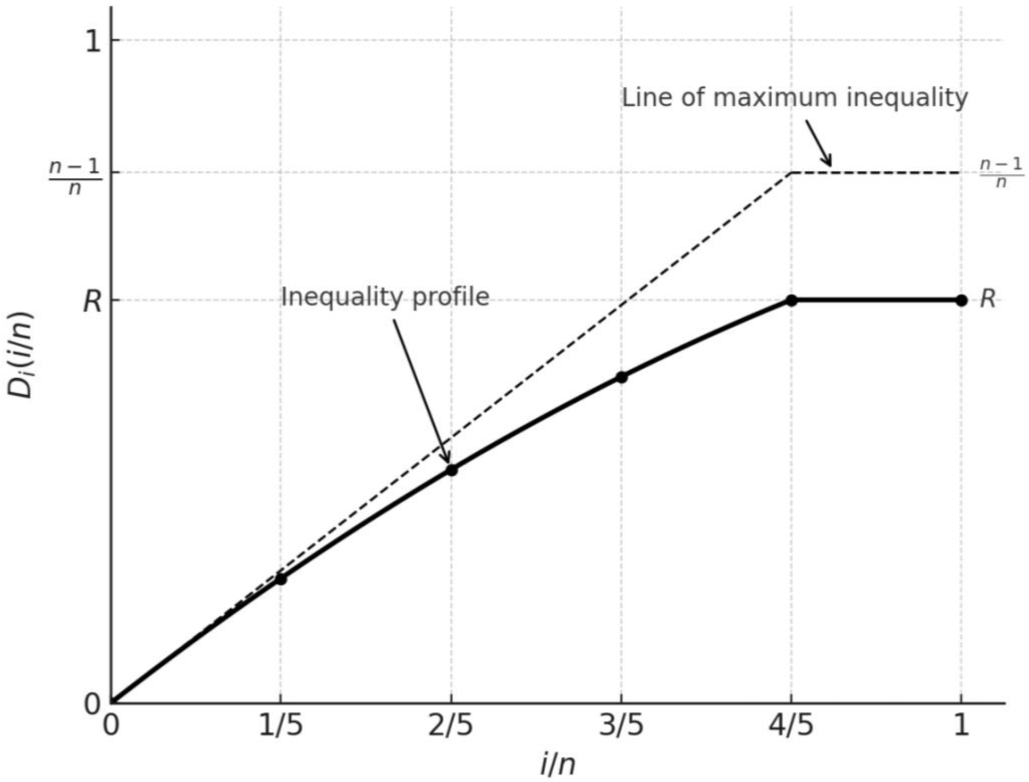

Figure 2: The inequality profile when *n* is 'large'

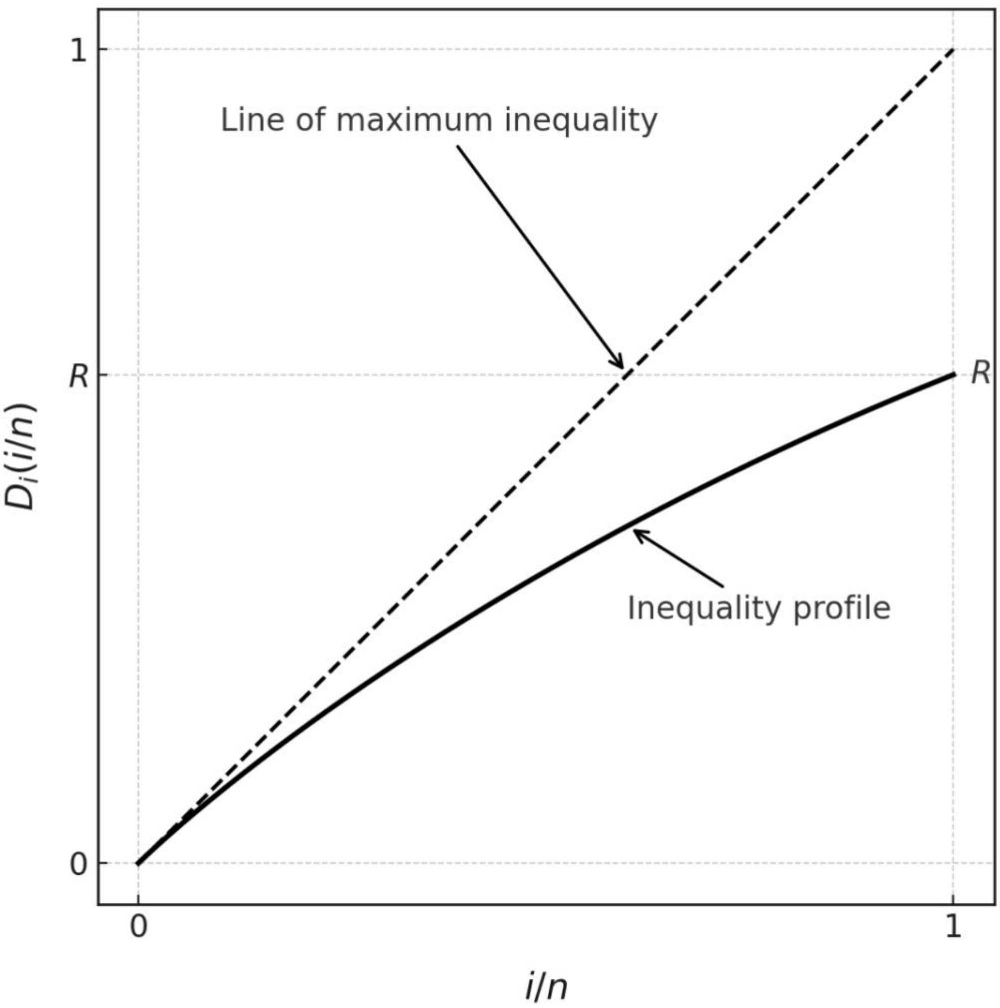

Figure 3: The relationship between *R* and the income share of the top 0.1%

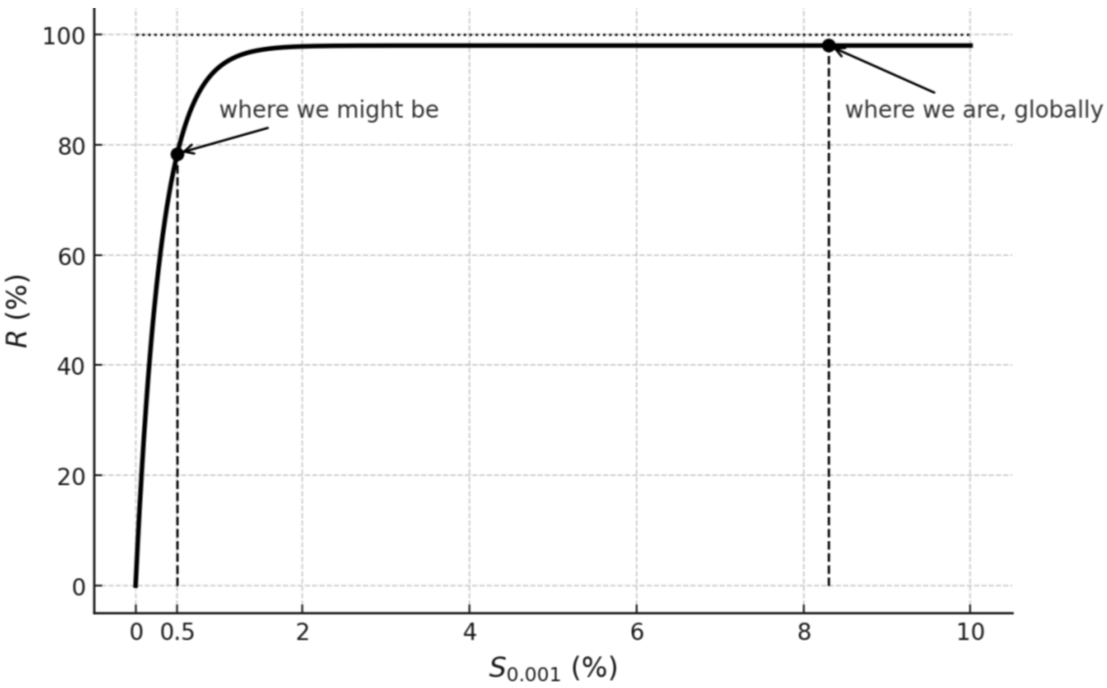

Appendix

Derivation of the Index *D* from the Inequality Profile

Figure A1 presents the inequality profile and line of maximum inequality in the standard unit square diagram for an assumed 5-person distribution. As noted in the text, the inequality measure *D* is the ratio of the area under the inequality profile to the area under the line of maximum inequality.

The inequality profile is a piece-wise linear curve obtained by joining the points $(1/5, D_1)$, $(2/5, D_2)$, $(3/5, D_3)$, $(4/5, D_4)$ and $(1, D_5)$, where (see the text):

$D_1 = (1/5x_5)(x_5 - x_1)$;

$D_2 = (1/5x_5)\{(x_5 - x_1) + (x_5 - x_2)\}$;

$D_3 = (1/5x_5)\{(x_5 - x_1) + (x_5 - x_2) + (x_5 - x_3)\}$;

$D_4 = (1/5x_5)\{(x_5 - x_1) + (x_5 - x_2) + (x_5 - x_3) + (x_5 - x_4)\}$; and

$D_5 = (1/5x_5)\{(x_5 - x_1) + (x_5 - x_2) + (x_5 - x_3) + (x_5 - x_4) + (x_5 - x_5)\}$.

The area under the inequality profile, call it *A*, is the sum of the areas of a number of triangles and rectangles, labelled a, b, c, d, e, f, g and h in Figure A1. Noting that the area of a triangle is (1/2)(base)(altitude) while the area of a rectangle is (base)(height), these various areas can be evaluated as follows, after defining $D_0 \equiv 0$:

Area a = $(1/2)(1/5)D_1 = (1/2)(1/5)(D_0 + D_1)$;

Area (b+c) = $(1/5)D_1 + (1/2)(1/5)(D_2 - D_1) = (1/2)(1/5)(D_1 + D_2)$; and, similarly,

Area (d+e) = $(1/2)(1/5)(D_2 + D_3)$;

Area (f+g) = $(1/2)(1/5)(D_3 + D_4)$; and

Area h = $(1/5)D_5 = (1/2)(1/5)(D_4 + D_5)$ since $D_5 = D_4$.

In general, and switching over to an *n*-dimensional income distribution,

Area A = $(1/2n)[(D_0 + D_1) + (D_1 + D_2) + ... + (D_{n-1} + D_n)]$, or, noting that $D_0 \equiv 0$,

Area A = $(1/2n)[2(D_1 + D_2 + ... + D_{n-1}) + D_n]$

= $(1/2n)[2(D_1 + D_2 + ... + D_{n-1} + D_n) - D_n]$, that is,

(A1) Area A = $(1/n)[(D_1 + ... + D_n) - D_n/2]$.

Recalling the definitions of the D_i, it should be clear that

$$D_1 + D_2 + ... + D_{n-1} + D_n =$$

$$(1/nx_n)[(x_n - x_1) + \{(x_n - x_1) + (x_n - x_2)\} + ...\{(x_n - x_1) + ... + (x_n - x_n)\}] =$$

$$(1/nx_n)[(1+2+...+n)x_n - \{nx_1 + (n-1)x_2 + ... + x_n\}], \text{or}$$

(A2) $D_1 + D_2 + ... + D_{n-1} + D_n = (1/nx_n)[\{n(n+1)/2\}x_n - \sum_{i=1}^{n}(n+1-1)x_i]$

where use has been made of the fact that the sum of the first n natural numbers is $n(n+1)/2$.

Now, if μ is the mean of the distribution, recall from Equation (1) in the text that the Gini coefficient of inequality can be written as

$G = (n+1)/n - (2/n^2\mu)\sum_{i=1}^{n}(n+1-i)x_i$, so that

$\sum_{i=1}^{n}(n+1-i)x_i = n(n+1)\mu/2 - n^2G\mu$, and substituting for $\sum_{i=1}^{n}(n+1-i)x_i$ and simplifying, (A2) can be written as

(A3) $D_1 + D_2 + ... + D_{n-1} + D_n =$

$$(1/2x_n)[(n+1)x_n - (n+1)\mu + nG\mu].$$

Substituting for $D_1 + D_2 + ... + D_{n-1} + D_n$ from (A3) into (A1) and recalling the definition of D_n enables us to write (A1) as:

Area $A = (1/2nx_n)[(n+1)x_n - (n+1)\mu + nG\mu] - (1/2n^2x_n)[(x_n - x_1) + ...(x_n - x_n)] =$

$$(1/2nx_n)[\{(n+1)x_n - (n+1)\mu + nG\mu\} - \{x_n - \mu\}] =$$

$$(1/2nx_n)[nx_n - n\mu(1-G)], \text{ or}$$

(A4) Area $A = (1/2)[1 - (\mu/x_n)(1-G)]$.

Finally, the area under the line of maximum inequality, call it A_{max}, is just the sum of the area of a right-angled triangle of base and height of $(n-1)/n$ each and a rectangle of base $1/n$ and height $(n-1)/n$, that is,

(A5) Area $A_{max} = (1/2)[(n-1)/n]^2 + (n-1)/n^2 = (1/2)[(n^2-1)/n^2]$.

The inequality measure D is just the ratio of the area A to the area A_{max}, so, in view of (A4) and (A5), we have:

$$D = \text{Area A/Area } A_{max} = \frac{(1/2)[1-(\mu/x_n)(1-G)]}{(1/2)[(n^2-1)/n^2]}, \text{ or}$$

$$D = \left(\frac{n^2}{n^2-1}\right)[1-(\mu/x_n)(1-G)] \text{ when } n \text{ is finite; and, as } n \to \infty,$$

$$D \to 1-(\mu/x_n)(1-G), \text{ as desired.}$$

Figure A1: The Inequality Profile for finite *n*, and the trapezoidal areas under the curve