IARIW 2021

Monday 23 – Friday 27 August

Statistisk sentralbyrå
Statistics Norway

Inequality Measurement for Bounded Variables

Inaki Permanyer

(Universitat Autonoma de Barcelona)

Suman Seth

(University of Leeds and University of Oxford)

Gaston Yalonetzky

(University of Leed)

Paper prepared for the 36th IARIW Virtual General Conference
August 23-27, 2021
Session 20: Inequality I
Time: Thursday, August 26, 2021 [14:00-16:00 CEST]

Inequality measurement for bounded variables*

Iñaki Permanyer† Suman Seth‡ Gaston Yalonetzky§

August 9, 2021

This is a preliminary draft; please do not cite without permission.

Abstract

Several variables of interest in social sciences are bounded. Hence many of them are represented in terms of either attainments or shortfalls. We jointly address two problems for inequality measurement posed by these potentially arbitrary choices. Firstly, inequality comparisons may not be robust when switching from attainment to shortfall representations. This is known as the consistency problem and the literature has provided several solutions for it. Secondly, maximum inequality with bounded variables is a predictable function of mean attainment (e.g. parabolic in the case of some well-known absolute inequality measures) characterised by low maximum levels when the mean approaches the bounds. This is known as the boundary problem and it compromises inequality comparisons of distributions with different means. In order to address these two concerns jointly, we propose a new inequality measurement framework for bounded variables, comprising a new class of normalised inequality indices that capture the extent to which maximum inequality is realised. The new inequality measurement framework is both consistent and immune to the boundary problem. We illustrate the efficacy of our approach through an example analysing cross-country international inequality in some well-known education and health indicators. Our approach show how the understanding of inequality using our proposed approach may produce a starkly different picture from the same using traditional inequality measurement frameworks.

Keywords: Inequality measurement, bounded variables, boundary problem, consistency, Kuznets curves.

JEL Codes: D63, I31, O47, O57.

*We would like to thank Florent Bresson, Ohiana Aristondo, Casilda Lasso de la Vega, Gordon Anderson, Jacques Silber, Stephan Klasen, Shatakshee Dhondge, and Conchita D'Ambrosio for very helpful comments and suggestions. We are also grateful to the participants of the 14th Meeting of the Society for Social Choice and Welfare, Seoul, South Korea; 13th conference of the Western Economic Association International, Catholic University, Santiago, Chile; the 13th Meeting of the Society for Social Choice and Welfare, Lund University, Sweden; the 7th Meeting of the Society for the Study of Economic Inequality, CUNY, USA; and the 2018 Royal Economic Society conference, University of Sussex, UK, for their valuable comments. Permanyer acknowledges funding from the European Research Council (ERC-2014-StG-637768, EQUALIZE project), the CERCA Programme from Generalitat de Catalunya and from the Spanish Ministry of Economy and Competitiveness "Ramón y Cajal" Research Grant Program (RYC-2013-14196). All remaining errors are our own.

†Centre d'Estudis Demografics, Universitat Autonoma de Barcelona. Email: ipermanyer@ced.uab.es.

‡Economics Division, Leeds University Business School, University of Leeds and Oxford Poverty & Human Development Initiative (OPHI), University of Oxford, Oxford, UK. Email: S.Seth@leeds.ac.uk.

§Leeds University Business School, University of Leeds. Email: G.Yalonetzky@leeds.ac.uk.

1 Introduction

In his seminal contribution, Atkinson (1970) set the foundations of inequality measurement as we know it. After five decades, the contributions to this burgeoning field of research have expanded in multiple directions, and 'inequality' can arguably be considered one of the most hotly debated topics in an increasingly globalised world, as witnessed by the popularity of several recent books on the subject (e.g. Piketty, 2015; Bourguignon, 2017; Atkinson, 2018; Milanovic, 2018). Moreover, the interest in inequality has gone well beyond the study of monetary distributions. Nowadays scholars and policy-makers alike are particularly interested in studying the distribution of the non-pecuniary dimensions of well-being (e.g. health and education outcomes, or the post-2015 Sustainable Development Goals [SDG] agenda).

Bounded variables abound. Unlike monetary variables such as income or consumption expenditure, the majority of non-pecuniary aspects of wellbeing are gauged by variables that cannot take indefinitely large values like literacy rates or infant mortality rates. This seemingly unimportant and technical point has key implications as regards the way in which we measure and interpret 'inequality' in the corresponding distributions. The measurement of inequality for bounded variables poses specific challenges that, while acknowledged in the literature, have not been *jointly* addressed in a satisfactory manner thus far. In this paper, we propose an approach to inequality measurement with bounded variables that addresses two key problems simultaneously.

The first problem that researchers encounter when studying inequality for bounded variables is the 'consistency problem'. When a variable is bounded it is a priori possible to focus either on the distribution of achievements or the corresponding distribution of shortfalls with respect to the upper bound.[1] As highlighted by Micklewright and Stewart (1999), Kenny (2004),Clarke et al. (2002), Erreygers (2009), Lambert and Zheng (2011), Lasso de la Vega and Aristondo (2012),Bosmans (2016) and many others, many inequality measures (especially popular relative measures like the Gini coefficient) fail to rank distributions consistently when measurement is switched from attainments to shortfall representations. Far from being a mere academic curiosity, the consistency problem poses several practical challenges to the study of inequality for bounded variables by precluding recourse to popular tools like the Lorenz curve. Fortunately, a battery of satisfactory solutions enabling the consistent measurement of inequality with bounded variables has been proposed, e.g. using absolute inequality measures (Erreygers, 2009; Lambert and Zheng, 2011), indices based on both representations (Lasso de la Vega and Aristondo, 2012), or using pairs of weakly consistent indices (Bosmans, 2016).

The second limitation is the 'boundary' problem (in some contexts also known as 'floor-' and 'ceiling-effect' problems). Whenever a variable is bounded, one can observe a clustering of the distribution as its mean converges towards any of its bounds. In these situations, the corresponding inequality levels mechanically go to zero, simply because there is no room left for further variation. This problem persists even when the consistency problem is solved using any of the aforementioned solutions. For instance, in the case of several absolute inequality measures and the absolute Lorenz curve, as the mean of the distribution increases from the lower bound to the upper bound of the distribution, the level of maximum feasible inequality first increases and then

[1] For instance, improvements in the coverage of public health plans could be assessed via the percentage of vaccinated children (an achievement indicator) or through the percentage of unvaccinated children (a shortfall indicator).

decreases; making the maximum feasible inequality a parabolic function of the mean. In these circumstances, it is not clear whether studying inequality with a bounded variable can provide new insights above and beyond what we already know from studying the values of the mean alone.[2] More generally, for any given consistent inequality measure, maximum feasible inequality becomes a predictable function of mean attainment making it difficult to disentangle 'mechanical' changes from inequality dynamics due to social phenomena when comparing distributions with different means.

In order to address these issues, we propose a new class of inequality indices, the so-called class of 'normalized inequality measures'. This class is obtained by combining a seemingly weak normalization axiom with the strong consistency requirement (Bosmans, 2016) and other minimally desirable properties (e.g. the transfers principle). As a result, we obtain new inequality measures quantifying observed inequality levels as a proportion of the maximum inequality levels that could be attained with the same index evaluated at an hypothetical distribution with the same mean as the observed distribution. The new measures satisfy the basic requirements of inequality measurement and are neither affected by the boundary effects nor by the consistency problems.

After defining our normalised inequality measures we illustrate how they perform empirically and compare them vis-à-vis their absolute and relative counterparts. We study the evolution of international inequality across countries in three education indicators from 1950 to 2010 and two health indicators from 1950 to 2015. Our results show that our newly proposed normalised inequality indices produce significantly different trend in inequality that the currently existing measures.

The rest of the paper is organised as follows. In section 2, we introduce the basic inequality measurement framework including notation and basic definitions. In sections 3 and 4 we examine, respectively, the challenges of consistency and boundary effects when measuring inequality with bounded variables. Section 5 presents the class of normalised inequality measures that provide consistent inequality comparisons for bounded variables free of the boundary problem. Section 6 presents the empirical illustration and section 7 concludes.

2 Basic framework for measuring inequality with bounded variables

Let us denote the set of all real numbers, all non-negative rational numbers and all positive natural numbers by $\mathbb{R}$, $\mathbb{Q}$ and $\mathbb{N}$, respectively. The non-negative and strictly positive counterparts of real numbers are denoted by $\mathbb{R}_+$ and $\mathbb{R}_{++}$, respectively. Suppose, there are n units of analysis (which could be people, households, municipalities or even countries, etc.) such that $n \in \mathbb{N} \setminus \{1\}$. Let $\mathbf{x} = (x_1, \ldots, x_n)$ be an *attainment distribution* of n units (or an n-dimensional *achievement vector*), where $x_i \in [0, U] \subseteq \mathbb{Q}$ represents unit i's cardinally measurable attainment bounded between a lower bound of zero and some fixed positive upper bound U.[3] We denote the set of all attainment distributions of size n with upper bound U by $\mathscr{X}_n \subseteq \mathbb{Q}_n$ and the set of all possible attainment distributions with upper bound U by $\mathscr{X} = \cup_n \mathscr{X}_n \subseteq \mathbb{Q}$.

[2] This point was already highlighted by Wagstaff (2005) in his study of the concentration index, and discussed by many others after him (e.g. Erreygers, 2009; Erreygers and Van Ourti, 2011; Wagstaff, 2009, 2011). In this paper we extrapolate some of these ideas to the context of inequality measurement for bounded variables.

[3] Our consideration for non-negative variables covers the vast majority of cases found in real-world applications.

Let us denote vectors comprising n ones and n zeros by $\mathbf{1}_n$ and $\mathbf{0}_n$, respectively. Consequently, for any $\lambda > 0$, $\lambda\mathbf{1}_n$ denotes the constant or egalitarian distribution where all n elements are equal to λ. We denote the arithmetic mean function evaluated at any $\mathbf{x} \in \mathscr{X}$ by $\mu(\mathbf{x})$. Additionally, for some $n \in \mathbb{N}$ and for some $U \in \mathbb{Q}$, let $\mathbb{G}_n = \{0, U/n, \ldots, (n-1)U/n, U\}$ denote a set of $n+1$ equally-spaced grid points between 0 and U. A distribution $\mathbf{x} \in \mathscr{X}_n$ is *bipolar* whenever for some $n' \in \mathbb{N}$ such that $n' < n$, n' units in $\mathbf{x}$ attain the value of U and the remaining $n - n'$ units attain the value of 0. Likewise, we refer to a distribution $\mathbf{x} \in \mathscr{X}_n$ as *almost-bipolar* whenever n' units in $\mathbf{x}$ attain the value of U, $n - n' - 1$ units in $\mathbf{x}$ attain the value of 0, and the leftover unit attains a value of $\varepsilon = [n\mu(\mathbf{x}) - n'U] \in (0, U)$. It will be clear afterwards that for a bipilar distribution $\mathbf{x} \in \mathscr{X}_n$, $\mu(\mathbf{x}) \in \mathbb{G}_n$.

An *inequality index* $I : \mathscr{X} \to \mathbb{R}_+$ is a continuous real-valued function and is expected to satisfy certain basic properties (Chakravarty, 2009). The first basic property, *anonymity*, requires that an inequality index should not depend on an eventual reordering of attainments across units. Formally, anonymity requires that $I(\mathbf{y}) = I(\mathbf{x})$ for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$ whenever $\mathbf{y} = \mathbf{xP}$, where $\mathbf{P}$ is a permutation matrix.[4] The second basic property, *transfer principle*, requires that a transfer from a richer to a poorer unit, without altering their relative positions, should decrease inequality (*progressive transfer*); whereas, alternatively, a transfer from a poorer to a richer unit should increase inequality (*regressive transfer*).[5] Formally, for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$, $I(\mathbf{y}) < I(\mathbf{x})$ whenever $\mathbf{y}$ is obtained from $\mathbf{x}$ by a progressive transfer and $I(\mathbf{y}) > I(\mathbf{x})$ whenever $\mathbf{y}$ is obtained from $\mathbf{x}$ by a regressive transfer.[6] The third basic property, *population principle*, requires that identical cloning of all units should leave inequality unaltered (thereby rendering populations with different sizes comparable). Formally, $I(\mathbf{y}) = I(\mathbf{x})$, whenever $\mathbf{y}$ is obtained from $\mathbf{x}$ by a *replication* for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}$.[7]

Finally, we refer to a real valued function $f : \mathscr{X} \to \mathbb{R}_+$ as *symmetric* whenever $f(\mathbf{x}) = f(\mathbf{xP})$ for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$, where $\mathbf{P}$ is a permutation matrix. We refer to a real valued function $f : \mathscr{X} \to \mathbb{R}_+$ as *strictly S-convex* if, for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$, $f(\mathbf{y}) < f(\mathbf{x})$ whenever $\mathbf{y}$ is obtained from $\mathbf{x}$ by a progressive transfer and $f(\mathbf{y}) > f(\mathbf{x})$ whenever $\mathbf{y}$ is obtained from $\mathbf{x}$ by a regressive transfer (Marshall and Olkin, 1979, p. 53-54).

3 The consistency problem

One issue that has received significant attention in the context of the measurement of inequality for bounded variables is the *consistency* problem. Variables with both lower and upper bounds (e.g. rates of mortality, literacy, access to basic facilities, etc.) can be represented by either their

[4] A *permutation matrix* is a square matrix with exactly one element in each row and column is equal to 1 and the rest of the elements are equal to zero.

[5] Technically, for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$, $\mathbf{y}$ is obtained from $\mathbf{x}$ by a *progressive transfer* whenever there are two units i, j and some $k > 0$ such that $y_i = x_i + k \leq x_j - k = y_j$ and $y_l = x_l$ for every $l \neq i, j$. Alternatively, $\mathbf{y}$ is obtained from $\mathbf{x}$ by a *regressive transfer* whenever there are two units i, j and some $k > 0$ such that $y_i + k = x_i \leq x_j = y_j - k$ and $y_l = x_l$ for every $l \neq i, j$.

[6] The bounded indicators discussed in this paper are not literally transferable (e.g., we do not suggest 'uneducating' highly educated individuals and transferring that education to less educated ones). Yet, one can compare the two hypothetical scenarios (pre- and post-"transfers") and still judge the latter exhibiting lower inequality than the former.

[7] Technically, $\mathbf{y} \in \mathscr{X}_m$ for some $m = \gamma n$, where $\gamma \in \mathbb{N} \setminus \{1\}$, is obtained from $\mathbf{x} \in \mathscr{X}_n$ by a *replication*, whenever $\mathbf{y} = (\mathbf{x}, \ldots, \mathbf{x})$, i.e. γ copies of $\mathbf{x}$ are repeated one after the other in $\mathbf{y}$.

distance from the lower bound (i.e. as attainments) or, alternatively, as their distance from the upper bound (i.e. as shortfalls). If $\mathbf{x} \in \mathscr{X}_n$ denotes the attainment distribution then we define the *shortfall distribution* associated with it as $\mathbf{x}^S = (x_1^S, \ldots, x_n^S) \in \mathscr{X}_n$ with $x_i^S = U - x_i$ representing unit i's shortfall with respect to the upper bound U. Since there is no a priori reason to prefer one representation (attainment or shortfall) over the other, the literature has proposed different versions of consistency properties that an inequality ordering should satisfy.

The first consistency property, *perfect complementarity*, requires that the value of the inequality index remains unaltered when we switch between attainment and shortfall representations of the same distribution, i.e., $I(\mathbf{x}) = I(\mathbf{x}^S)$ for any $\mathbf{x} \in \mathscr{X}$ (Erreygers, 2009). The second consistency property, *strong consistency*, requires that inequality measures should rank pairs of attainment distributions and their shortfall counterparts in a coherent manner. In other words, the inequality ranking should be robust to alternative representations of the variable., i.e., $I(\mathbf{x}) \leq I(\mathbf{y}) \Leftrightarrow I(\mathbf{x}^S) \leq I(\mathbf{y}^S)$ for any $\mathbf{x}, \mathbf{y} \in \mathscr{X}$ (Lambert and Zheng, 2011). Clearly, perfect complementarity implies strong consistency, but the reverse is not true. The third consistency property, *weak consistency*, proposed by Bosmans (2016), is predicted on the realisation that it is possible to find pairs of different inequality indices that produce consistent comparisons as long as one index I^A is used for the attainment distribution and another index $I^S = \phi(I^A)$ is used for the shortfall counterpart, where $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a strictly increasing function. The pair (I^A, I^S) is *jointly* weakly consistent if and only if $I^A(\mathbf{x}) \leq I^A(\mathbf{y}) \Leftrightarrow I^S(\mathbf{x}^S) \leq I^S(\mathbf{y}^S)$ for any $\mathbf{x}, \mathbf{y} \in \mathscr{X}$. For example, if $I^A(\mathbf{x})$ is the Gini coefficient evaluated on the attainment distribution $\mathbf{x}$, then $I^S(\mathbf{x}^S) = \mu(\mathbf{x}^S) I^A(\mathbf{x}^S)/[U - \mu(\mathbf{x}^S)]$ may be a potential candidate for consistently evaluating inequality for the shortfall distribution.

Different consistency properties restrict the set of permissible inequality measures. For instance, the advocates of strong consistency (Erreygers, 2009; Lambert and Zheng, 2011) suggest using absolute inequality indices (and related partial orderings). An inequality measure is *absolute* if the evaluation of inequality remains unchanged when all attainments are changed by the same amount, i.e., if it satisfies the *translation invariance* property, which requires that $I(\mathbf{y}) = I(\mathbf{x})$ for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$ whenever $\mathbf{y} = \mathbf{x} + \lambda \mathbf{1}_n$ for some $\lambda \in \mathbb{R}$.[8] Lasso de la Vega and Aristondo (2012), however, showed that strong consistency is satisfied by a wider class of inequality indices that stem from equally weighted generalised means of an inequality index evaluated at the attainment distribution and the same index evaluated at the corresponding shortfall distribution. Finally, Bosmans (2016) showed that relaxing strong consistency to weak consistency allows permissible measures to include even relative inequality indices. An inequality index is *relative* if the evaluation of inequality remains unchanged when all attainments are changed by the same proportion, i.e., if it satisfies the *scale invariance* property, which requires that $I(\mathbf{y}) = I(\mathbf{x})$ for some $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$ whenever $\mathbf{y} = \lambda \mathbf{x}$ for some $\lambda > 0$. Weak consistency certainly broadens the set of permissible inequality measures, but strong consistency may appear to be more amenable to practical applications.

[8] Chakravarty et al. (2015) also identify a family of absolute consistent inequality indices. Seth and Alkire (2017) adopt an absolute approach to consistently measure inequality among the poor in the counting poverty framework.

4 The boundary problem

By comparison to the consistency problem, the *boundary problem* has received little attention. Generally, the boundary problem stems from the existence of a predictable functional relation between mean attainment and maximum feasible inequality (for a given inequality index), which hinders comparisons between distributions of bounded variables with different means. Indeed, when means differ it can be hard to disentangle the change in inequality due to social phenomena (arguably the one we are interested in) from that due to the predicable relationship between mean attainment and maximum feasible inequality. For instance, in the case of many absolute inequality measures like the absolute Gini and the variance (which are consistent) when mean attainment converges toward any of the bounds maximum feasible inequality falls which prompts observed inequality to fall in consequence. In fact, with these absolute inequality measures the relationship between mean attainment and maximum feasible inequality is usually parabolic with a global maximum at $\mu(\mathbf{x}) = 0.5U$ (if the lower bound is equal to 0).

Figure 1: Change in the mean and different absolute inequality measures for cross-country Hepatitis-B immunization rates

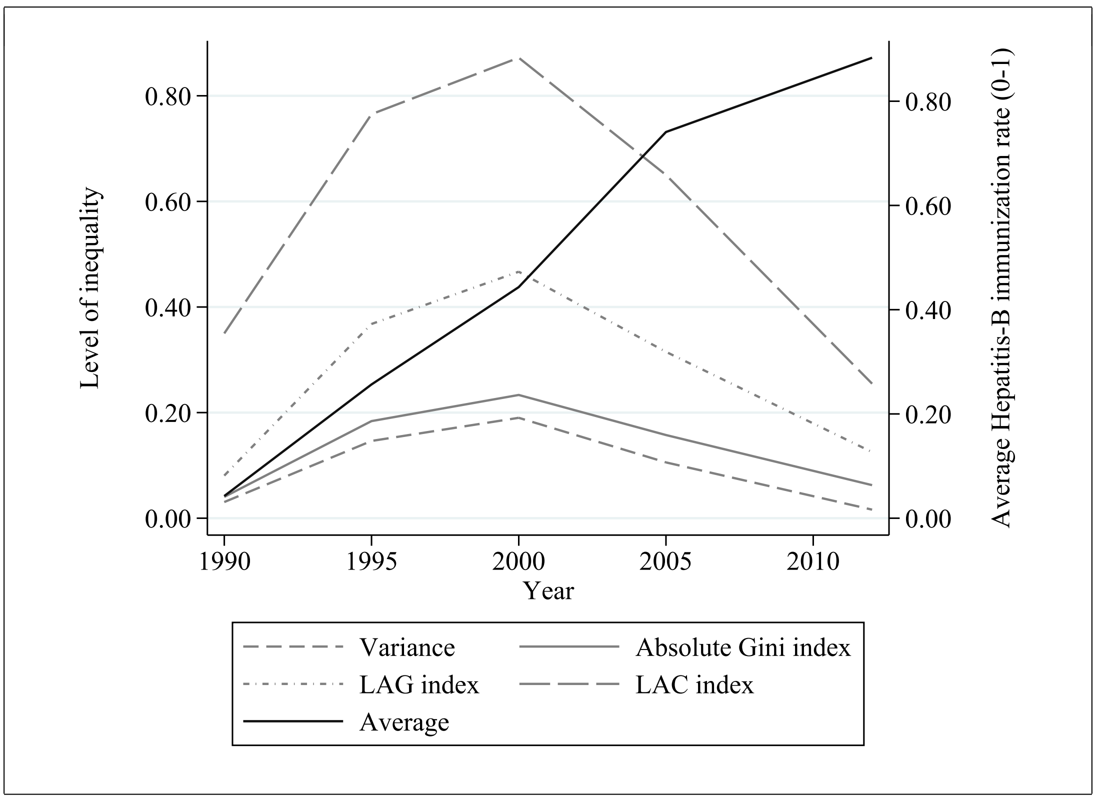

Source: Authors' own computations.

Notes: The formula for the variance is $V(\mathbf{x}) = (\sum_{i=1}^{n}[x_i - \mu(\mathbf{x})]^2)/n$. The formula for the absolute Gini index is $G_a(\mathbf{x}) = \mu(\mathbf{x})G_r(\mathbf{x})$ where G_r is the relative Gini index: $G_r(\mathbf{x}) = (\sum_{i=1}^{n}\sum_{j=1}^{n}|x_i - x_j|)/2\mu(\mathbf{x})n^2$. The LAG index is equal to $LAG(\mathbf{x}) = ([G_r(\mathbf{x})^{-1} + G_r(\mathbf{x}^S)^{-1}]/2)^{-1}$. The LAC index is equal to $LAC(\mathbf{x}) = ([CV(\mathbf{x})^{-1} + CV(\mathbf{x}^S)^{-1}]/2)^{-1}$, where $CV \equiv \sqrt{V}/\mu$ is the coefficient of variation.

In order to illustrate the empirical relevance of the boundary problem, figure 1 shows the change in mean Hepatitis-B immunization rates across countries as well as the changes in various measures of cross-country inequality between 1990 and 2012. We use four different consistent inequality measures: two absolute inequality indices, the Variance V and the absolute Gini coefficient G_a (relative Gini coefficient *times* the mean); and two inequality measures from the class of measures proposed by Lasso de la Vega and Aristondo (2012) (see note in figure 1 for details). As mean attainment increases from the lower bound to the upper bound, the values of all four inequality measures first increase and then decrease.

Thus, we observe a *Kuznets curve* relating mean attainment to each of the four consistent inequality measures over time. Is this type of relationship merely an empirical regularity or a mechanical artifact driven by the susceptibility of existing inequality indices (consistent or not) to the boundary problem?

In order to explore whether there is a predictable relationship between mean attainment and maximum possible inequality in the distribution of a bounded variable, first we identify whether such distribution reflecting maximum inequality for a given mean exists and, if so, what it looks like. Then we need to check whether inequality indices evaluated at such distributions change in predictable ways when the mean is different. Proposition 1 shows that a *maximum-inequality distribution* (MID), $\hat{\mathbf{x}} \in \mathscr{X}_n$, for a given distribution $\mathbf{x} \in \mathscr{X}_n$ with $\mu(\mathbf{x}) = \mu(\hat{\mathbf{x}})$ exists and how it looks like.

Proposition 1 For any $n \in \mathbb{N} \setminus \{1\}$ and for any $\mathbf{x} \in \mathscr{X}_n$, a maximum inequality distribution $\hat{\mathbf{x}} \in \mathscr{X}_n$, such that $\mu(\mathbf{x}) = \mu(\hat{\mathbf{x}})$, exists and is either *bipolar* or *almost-bipolar*.

Proof. See Appendix A1. ■

As it turns out, an MID is either bipolar or almost bipolar. Bipolar distributions consist of frequencies concentrated either at the lower bound or upper bound exclusively; whereas, almost bipolar distributions consists of all units with either the lower or upper bound value, except for one unit with interior value $\varepsilon \in (0,U)$. For example, assume $U = 1$, $n = 4$ and so $\mathbb{G} = \{0, 0.25, 0.5, 0.75, 1\}$. Consider first the distribution $\mathbf{x} = (0.1, 0.4, 0.7, 0.8)$ with $\mu(\mathbf{x}) = 0.5 \in \mathbb{G}$. In this case, the corresponding MID is $\hat{\mathbf{x}} = (0,0,1,1)$, which is bipolar, and clearly $\mu(\hat{\mathbf{x}}) = \mu(\mathbf{x}) = 0.5$. Now consider a second distribution $\mathbf{y} = (0.2, 0.4, 0.7, 0.9)$ with $\mu(\mathbf{y}) = 0.55 \notin \mathbb{G}$. The corresponding MID, in this case, is $\hat{\mathbf{y}} = (0, 0.2, 1, 1)$ with $\mu(\hat{\mathbf{y}}) = \mu(\mathbf{y}) = 0.55$ and $\varepsilon = 0.2 \in (0,1)$, but $\hat{\mathbf{y}}$ is almost-bipolar and no further regressive transfer is possible. Note that an MID is unique for a given mean and for a given n. In other words, for any two distributions $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$ such that $\mathbf{x} \neq \mathbf{y}$ but $\mu(\mathbf{x}) = \mu(\mathbf{y})$, it is always $\hat{\mathbf{x}} = \hat{\mathbf{y}}$.

Even though MIDs are hypothetical distributions unlikely to be observed in practice, they do represent the benchmark case of maximum inequality against which we can compare bounded distributions. An inequality evaluation of a distribution cannot be larger than its MID as long as an inequality index I satisfies anonymity and the transfer principle.

Corollary 1 For any $n \in \mathbb{N} \setminus \{1\}$ and for any $\mathbf{x} \in \mathscr{X}_n$, $I(\mathbf{x}) \leq I(\hat{\mathbf{x}})$ whenever I satisfies anonymity and the transfer principle, where $\hat{\mathbf{x}}$ is the MID for $\mathbf{x}$.

Figure 2: The relationship between the mean and the maximum inequality of a bounded variable

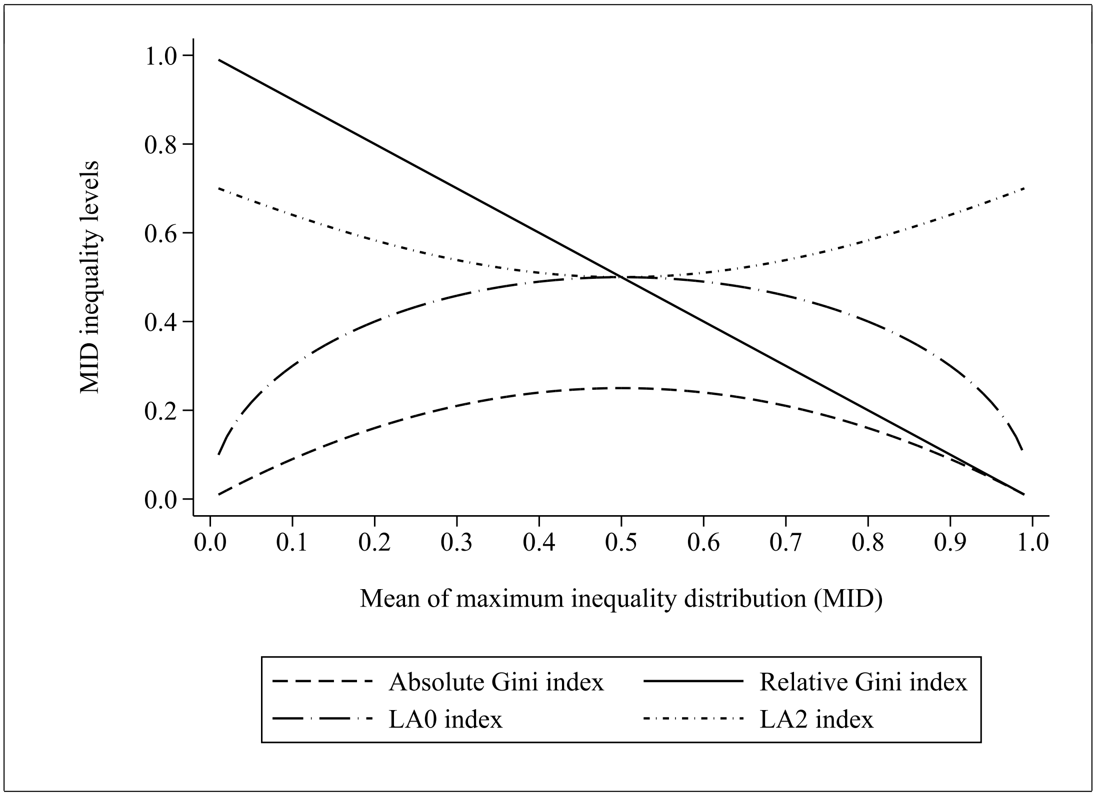

Source: Authors' own computations.
Notes: The figure is based on $U = 1$ and bipolar distributions. $G_a(\hat{\mathbf{x}}) = \mu(\hat{\mathbf{x}})[1-\mu(\hat{\mathbf{x}})]$ and $G(\hat{\mathbf{x}}) = 1-\mu(\hat{\mathbf{x}})$. LA0 and LA2 are indices from the class proposed by Lasso de la Vega and Aristondo (2012). LA0 is the geometric mean of the (relative) Gini index evaluated at $\mathbf{x}$ and the same index evaluated at $\mathbf{x}^S$, whereas LA2 is the Euclidean mean of those two indices. Thus $LA0(\hat{\mathbf{x}}) = \sqrt{\mu(\hat{\mathbf{x}})[1-\mu(\hat{\mathbf{x}})]}$ and $LA2(\hat{\mathbf{x}}) = \sqrt{[\mu(\hat{\mathbf{x}})^2 + (1-\mu(\hat{\mathbf{x}}))^2]/2}$.

Proof. Given that $\hat{\mathbf{x}}$ is obtained from $\mathbf{x}$ by one or more regressive transfers, $I(\mathbf{x}) < I(\hat{\mathbf{x}})$ by the transfer principle. Whenever $\mathbf{x} = \hat{\mathbf{x}}$, $I(\mathbf{x}) = I(\hat{\mathbf{x}})$ by anonymity. ■

Figure 2 shows the relationship between mean and inequality for bipolar MIDs using four different inequality measures that satisfy some form of consistency. We assume that $U = 1$ and $0 < \mu(\mathbf{x}) < 1$ in order to allow the existence of inequality. Moreover, $\mu(\mathbf{x}) \in \mathbb{G}$ because the constructed MIDs are bipolar. We choose the strongly-consistent absolute Gini index (Lambert and Zheng, 2011); the weakly-consistent relative Gini index (Bosmans, 2016); and two strongly-consistent indices from the class proposed by Lasso de la Vega and Aristondo (2012): the equally weighted geometric mean (LA0) and Euclidean mean (LA2), respectively, of the relative Gini indices evaluated at attainment and corresponding shortfall distributions. As $\mu(\mathbf{x})$ increases in figure 2, we move from one MID to another and the inequality measures react differently to the changes in the MID. For example, the absolute Gini index and the LA0 index first increase, reach their maximum values around $\mu(\mathbf{x}) = 0.5$ and then decrease. By contrast, the LA2 index first decreases, reaches its minimum value around $\mu(\mathbf{x}) = 0.5$ and then increases. Finally, the relative Gini index decreases monotonically as the mean increases.

Evidently, inequality comparisons of distributions with *different means* using existing inequality indices suffer from the boundary problem, thereby not being sufficiently meaningful unless we control for changes in maximum possible inequality as a function of the mean. In order to control for the effect of differing means, we propose a simple approach. For a given distribution, we suggest comparing the observed inequality level against that in the corresponding MID. Our proposal involves using a 'relative-like' inequality measure that assesses how large an observed inequality level is with respect to the extent of inequality in the corresponding counterfactual distribution that serves as a reference point.

5 A family of normalised inequality indices

Our proposal to solve the boundary problem requires measuring observed inequality as a proportion of the maximum attainable with the same mean attainment, i.e. evaluated at the respective MID. The proposal restores comparability across distributions with different means and (strongly) consistently across alternative representations. The crucial property imposed to counter the boundary problem is the *maximality principle*:[9]

Maximality Principle For any $\mathbf{x} \in \mathscr{X}_n$, $I(\mathbf{x}) = 1$ whenever $\mathbf{x} = \hat{\mathbf{x}}$.

As we will show, the maximality principle ensures that inequality is measured as a proportion of the maximum level reachable given a mean attainment so that any distribution different from the corresponding MID obtains an inequality value strictly below 1.

We also introduce the standard *equality principle*, which requires that $I(\mathbf{x}) = 0$ whenever $\mathbf{x} = \lambda \mathbf{1}_n$ for any $\mathbf{x} \in \mathscr{X}_n$ and $\lambda \geq 0$. This property ensures that inequality is minimal and equal to zero whenever all units feature exactly the same value for the indicator, i.e. $x_1 = x_2 = \cdots = x_n$. From now onwards, we refer to the family of inequality indices satisfying the maximality principle and the equality principle in addition to other desirable properties as *normalised inequality indices*, which is characterized in theorem 1.

Theorem 1 For any $n \in \mathbb{N} \setminus \{1\}$ and any $\mathbf{x} \in \mathscr{X}_n$, an inequality index I satisfies anonymity, the transfer principle, the equality principle, the maximality principle and strong consistency whenever

$$I(\mathbf{x}) = \begin{cases} \dfrac{f(\mathbf{x}) - f(\bar{\mathbf{x}})}{f(\hat{\mathbf{x}}) - f(\bar{\mathbf{x}})} & \text{for } \mathbf{x} \neq \bar{\mathbf{x}} \\ 0 & \text{for } \mathbf{x} = \bar{\mathbf{x}} \end{cases},$$

where $\bar{\mathbf{x}} = \mu(\mathbf{x})\mathbf{1}_n$, $\hat{\mathbf{x}}$ is the MID for $\mathbf{x}$, and $f : \mathscr{X}_n \to \mathbb{R}_{++}$ is a symmetric and strictly S-convex function with $f(\mathbf{x}^S) - f\left(\bar{\mathbf{x}}^S\right) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ for some $m, c \in \mathbb{R}_{++}$.

Proof. See Appendix A2. ■

[9] MIDs can also be identified with relative inequality measures applied to variables without upper bounds and negative values. In such cases the MID is characterised by one unit having positive income and everyone else featuring zero income. The values of relative inequality indices evaluated at these MIDs often depend on the population size. Indeed that is the case of the relative Gini index which is normalised to attain the value of $1 - \frac{1}{n}$ when evaluated at a MID. Otherwise normalised relative inequality indices are not popular even though they can be constructed. For normalised versions of members of the generalised entropy class see Shorrocks (1980, footnote 7, p. 623).

Theorem 1 states that an inequality measure that is strongly consistent *and* normalised between 0 and 1 according to the equality principle and the maximum inequality axiom needs to be a particular type of weakly consistent inequality index (since $f(\mathbf{x}) - f(\bar{\mathbf{x}})$ satisfies all the minimum properties of an inequality index, namely anonymity, transfers principle and the equality principle, and $f(\mathbf{x}^S) - f\left(\bar{\mathbf{x}}^S\right) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ implies weak consistency though the reverse is not true) divided by its maximum value which, in turn, depends on the MID associated with any particular distribution (meaning that $\mu(\mathbf{x}) = \mu(\hat{\mathbf{x}})$).

Crucially, note that aside from $f(\mathbf{x}^S) - f\left(\bar{\mathbf{x}}^S\right) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ no other particular requirement is imposed on the inequality index in the numerator and denominator (i.e. I) in terms of its sensitivity to rescaling, translations or other similar transformations defining the different classes of relative, absolute, intermediate, super-relative or super-absolute measures.[10] This means that several members of the aforementioned classes are suitable for being in the numerator and denominator of I. For instance, in the case of absolute indices like the variance, the absolute Gini or the members of the class proposed by Chakravarty et al. (2015), $f(\mathbf{x}^S) - f\left(\bar{\mathbf{x}}^S\right) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ is fulfilled with $C = m = 1$ (noting that in all these cases $f\left(\bar{\mathbf{x}}^S\right) = f(\bar{\mathbf{x}}) = 0$). The same goes for all members of the class proposed by Lasso de la Vega and Aristondo (2012) (i.e. $C = m = 1$). Likewise, both the (relative) Gini index and the coefficient of variation satisfy the same equality with $m = 1$ and $C = \frac{\mu(\mathbf{x})}{\mu(\mathbf{x}^S)}$. In summary, the class of normalised inequality indices admits myriad functional forms in the numerator and denominator.

The formulas of the normalised Gini index and the normalized variance appear, respectively, in expressions 1 and 2:

$$G_U^*(\mathbf{x}) = \begin{cases} \dfrac{G_a(\mathbf{x})U}{\mu(\mathbf{x})(U - \mu(\mathbf{x}))} & \text{if } \mathbf{x} \neq \bar{\mathbf{x}} \\ 0 & \text{if } \mathbf{x} = \bar{\mathbf{x}} \end{cases}, \tag{1}$$

$$V_U^*(\mathbf{x}) = \begin{cases} \dfrac{V(\mathbf{x})}{\mu(\mathbf{x})(U - \mu(\mathbf{x}))} & \text{if } \mathbf{x} \neq \bar{\mathbf{x}} \\ 0 & \text{if } \mathbf{x} = \bar{\mathbf{x}} \end{cases}. \tag{2}$$

6 Empirical illustration: Evolution of cross-country inequality in education and health

In this section, we present some empirical illustrations to reflect the efficacy of our proposal, by studying the evolution of cross-country inequality in three education indicators and two health indicators since 1950. To examine the evolution of inequality in education, we look at the *share of total adult population with at least some primary education*, the *share of total adult population with at least some secondary education* and the *share of total adult population with at least some tertiary education*. To examine inequality in health, we select the under five survival rate and the infant survival rate, where survival rates are the complements of the respective mortality rates. The education data were obtained from the Barro-Lee dataset; whereas, the data on mortality rates

[10] See Bosmans (2016) for a concise and comprehensive typology.

were obtained from the United Nations' Department of Economic and Social Welfare website.[11] The three education indicators reflect rates and so they are bounded between zero and one. For the health indicators, we obtain the survival rates by subtracting the mortality rates from 1,000 and then normalising the shortfall by 1,000. Therefore, the survival rates also lie between zero and one.

Figure 3: Changes in the means of selected education indicators and health indicators since 1950

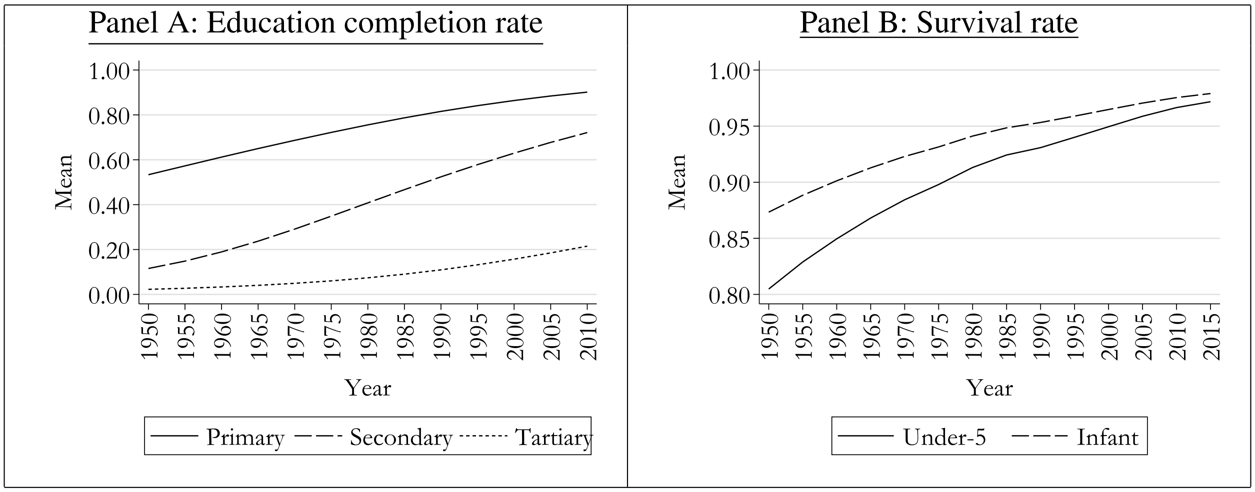

Source: Authors' own computations from the datasets used.
Notes: We have used the data for 133 countries for the three education indicators and the data for 201 countries for the two health indicators.

In Figure 3, we present the change in mean attainments for the education indicators between 1950 and 2010 and for the health indicators between 1950 and 2015, for every five year period.[12] Global average of all five indicators of education and health show steady improvements since 1950. However, the means lie at different intervals for different indicators. Between 1950 and 2010, the mean for the primary education completion rate indicator increases from 0.53 to 0.92, the mean for the secondary education completion rate indicator increases from 0.12 to 0.76, and the tertiary education completion rate indicator increases from 0.02 to 0.25. Note that the mean of the primary education completion rate indicator lies above 0.50 throughout and that of the tertiary education completion rate indicator lies below 0.50 throughout, but the same for the secondary education completion rate indicator lies below 0.50 before 1990 and then it surpasses 0.50. For the two health indicators, however, the mean of the under-5 survival rate indicator increases from 0.80 in 1950 to 0.97 in 2015 and the mean of the infant survival rate indicator increases from 0.87 in 1950 to 0.98 in 2015. Note that the means of both health indicators lie well above 0.50 throughout.

In Figure 4, we present absolute Gini indices for all five indicators and their relationships to the respective mean attainments over time. Panels A1 and A2 in the figure present the trends in absolute Gini since 1950. Interesting trends emerge when we look at different indicators. Absolute inequality across countries for the primary education completion rate indicator falls throughout, which is compatible to the fact that the mean attainment in the indicator has been more than 0.5

[11] Source of the education data: http://www.barrolee.com/. Source of health data: https://population.un.org/wpp/Download/Standard/Mortality/.

[12] The mean attainment does not include population weights. Each country, irrespective of its size, is considered as an unit with equal importance.

Figure 4: Change in cross-country absolute Gini indices for education and health indicators since 1950

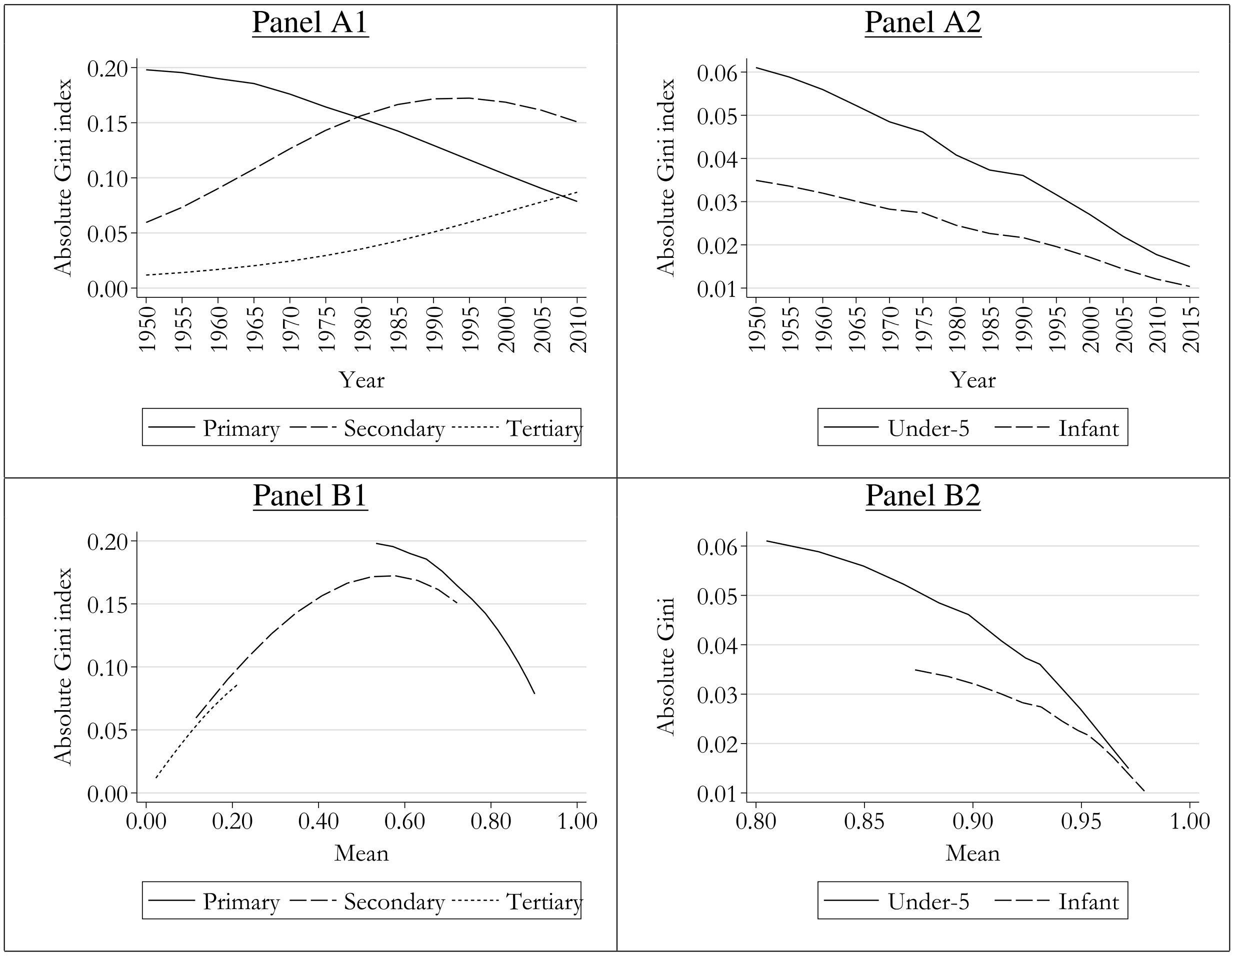

Source: Authors' own computations from the datasets used.
Notes: We have used the data for 133 countries for the three education indicators and the data for 201 countries for the two health indicators. The absolute Gini index is $G_a(\mathbf{x}) = \mu(\mathbf{x})G_r(\mathbf{x})$ where G_r is the relative Gini index: $G_r(\mathbf{x}) = (\sum_{i=1}^{n}\sum_{j=1}^{n}|x_i - x_j|)/2\mu(\mathbf{x})n^2$.

since 1950. Similarly, absolute inequality across countries for the tertiary education completion rate indicator rises throughout and is compatible to the fact that the mean attainment in the indicator has been always less than 0.5 since 1950. Interesting pattern is observed for the secondary education completion rate indicator, where inequality rises and then falls after around 1990, which is also compatible with the fact that the mean of the indicator surpasses 0.50 around that year. The trends for absolute Gini indices (Panel A2) for the two health indicators are the same as that of the primary education completion rate indicator as the means lie above 0.50 for the entire period. The relationships between mean attainments and absolute Gini indices are presented in the Panels B1 and B2, where each horizontal axis represent the mean attainment and the vertical axis represents the absolute Gini index. The relationship between the mean and absolute Gini index appears to be quite mechanical for our illustration.

In Figure 5, we present the evolution of relative inequality with the relative Gini index. Unlike the absolute Gini indices in Figure 4, the patterns for relative Gini indices are uniform across

Figure 5: Change in cross-country relative Gini indices for education and health indicators since 1950

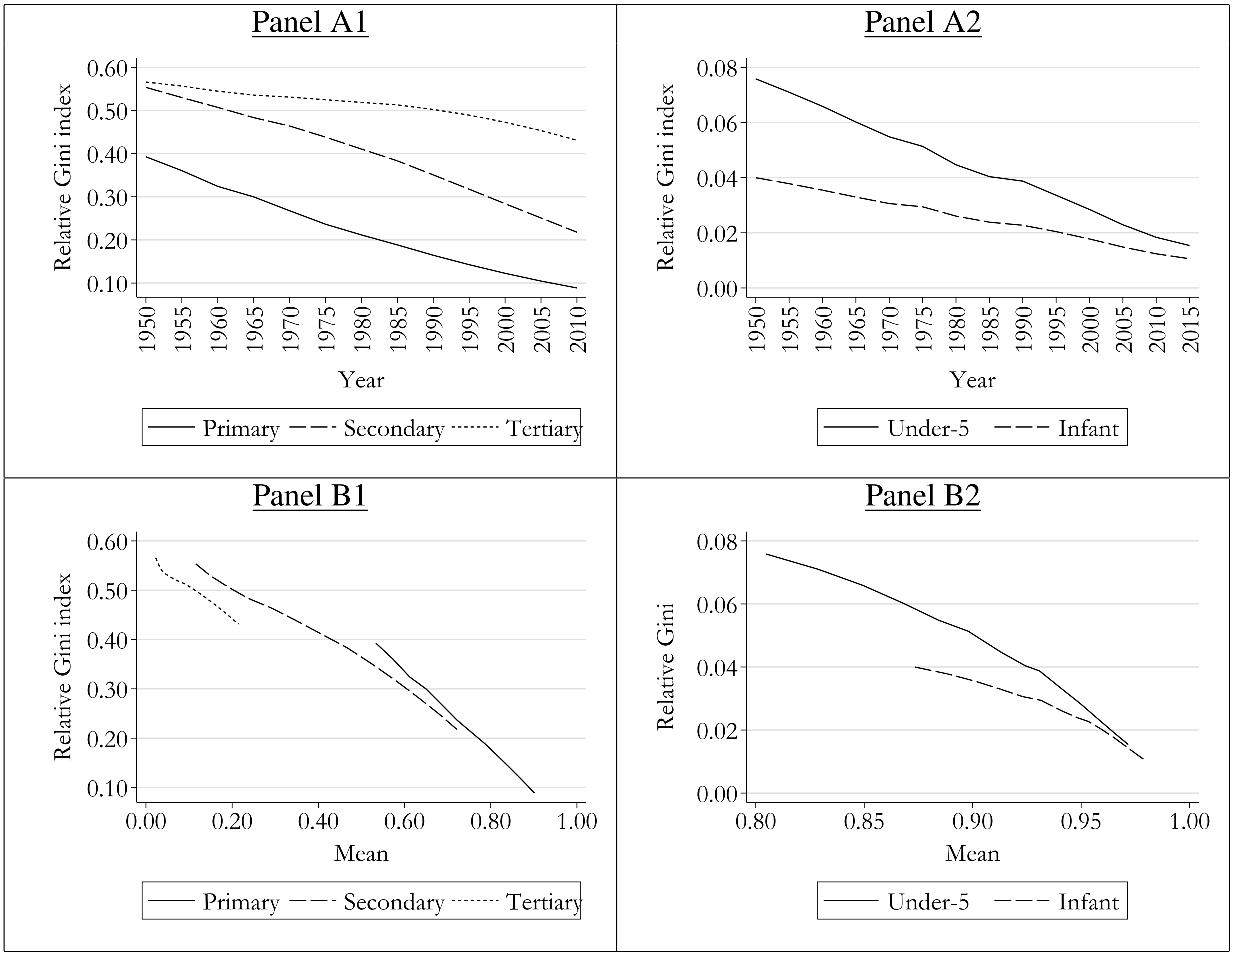

Source: Authors' own computations from the datasets used.

Notes: We have used the data for 133 countries for the three education indicators and the data for 201 countries for the two health indicators. The relative Gini index is formulated as $G_r(\mathbf{x}) = (\sum_{i=1}^{n}\sum_{j=1}^{n}|x_i - x_j|)/2\mu(\mathbf{x})n^2$.

all education and health indicators. As mean attainments increase steadily for all five indicators, relative Gini indices appear to fall throughout, which is evident from Panels A1 and A2 of the figure. The relationships between the mean attainments and relative Gini indices are depicted in Panels B1 and B2 of the figure. Both absolute and relative Gini indices follow the pattern that we discussed in Section 4.

Finally, in Figure 6, we present the trends in normalised Gini indices, where we normalise the absolute Gini index by the maximum attainable Gini index for a given level of mean attainment for a bounded variable. Again, we observe, in Panels A1 and A2 of Figure 6, that different trends emerge. Let us start with the primary education completion rate indicator. Unlike both absolute and relative Gini indices, the normalised Gini indices register an upward trend. Therefore, after we control for the boundary problem, cross-country inequality for the primary education completion rate indicator does not appear to have gone down, as reflected by both absolute and relative Gini indices. The normalised Gini indices for the secondary education completion rate indicator follow the same trend between 1950 and 2010, but the normalised Gini indices for the tertiary education

Figure 6: Change in cross-country normalised Gini indices for education and health indicators since 1950

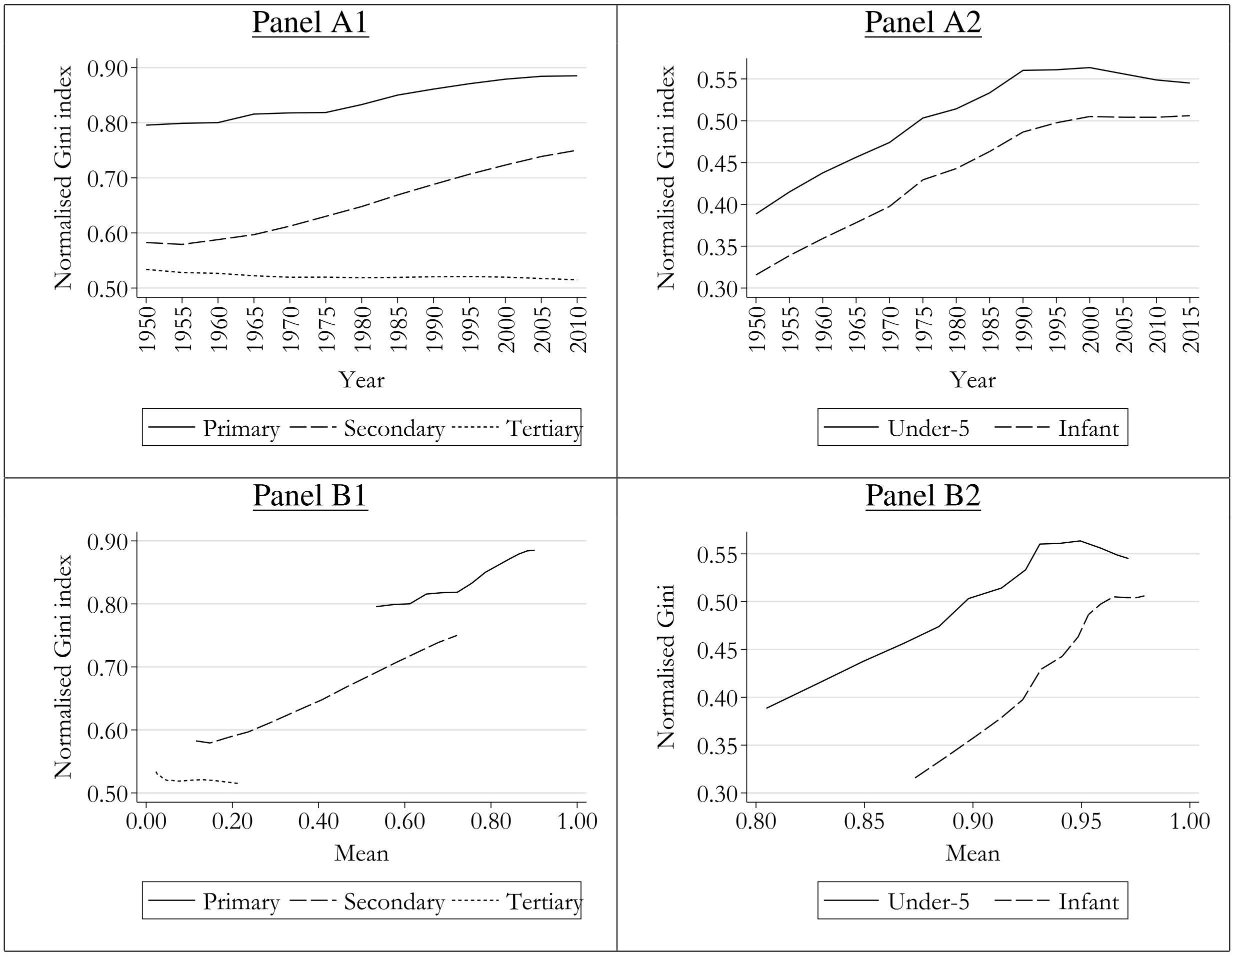

Source: Authors' own computations from the datasets used.
Notes: We have used the data for 133 countries for the three education indicators and the data for 201 countries for the two health indicators. The normalised Gini index is $G_U^*(\mathbf{x}) = G_a(\mathbf{x})/[\mu(\mathbf{x})(1-\mu(\mathbf{x}))]$ where G_a is the absolute Gini index.

completion rate indicator linger between 0.51 and 0.54 for the entire period. Both health indicators follow the trend for the primary education completion rate indicator for both absolute Gini indices and relative Gini indices, but they do not follow the same trend for the normalised Gini indices for the entire period. The normalised Gini indices for the under-5 survival rate indicator increase until year 2000 and then fall. Similarly, the normalised Gini indices for the infant survival rate indicator also increase until year 2000, but then they stabilise. Therefore, what we observe from our analysis based on the illustration is that normalised Gini indices may produce very different trends for inequality in practice than traditional absolute and relative indices.

7 Summary and concluding remarks

The use of traditional income inequality measures to study the variability of bounded variables poses several problems. On the one hand, the bounded nature of the variables generate a predictable functional relationship between the mean of a distribution and its maximum inequality levels (e.g. in the case of many indices, chiefly absolute ones, when the mean approaches either the upper or the lower bound, inequality mechanically goes to zero). On the other hand, several measures of inequality, e.g. all the relative ones, fail to rank distribution pairs consistently between the alternatives of attainments and shortfall representations. We proposed a new approach to inequality measurement aimed at solving both problems, consistency and boundary effects, *simultaneously*. Basically, the normalised inequality measures compare observed inequality levels against the maximum inequality level achievable with the same measure across all hypothetical distributions having the same mean. The normalised inequality indices are strongly consistent (Bosmans, 2016) and eliminate any mechanical relationship between mean attainment and maximum inequality brought about by the boundary problem.

To illustrate our approach we have investigated the evolution of international inequality in education and health (i.e. between-country inequality in life expectancy) from 1950 to the recent periods using the normalised inequality measures.

While our proposed inequality measurement framework successfully addresses the consistency problem and the boundary effect, there remain other measurement challenges in the setting of bounded variables For example, Lasso de la Vega and Aristondo (2012) provide conditions whose fulfillment guarantees robustness of inequality comparisons to changes in the upper bound. While admittedly this problem is not that serious when bounds are neither arbitrary nor expected to change across time and space (e.g. in the case of indicators expressed as percentage ratios), it is nonetheless worth exploring how our proposed inequality measurement framework could also accommodate this potential concern. This challenge is left for future research. Likewise, future research could explore partial orderings respecting the consistency and normalisation properties that were combined to generate the normalised inequality indices.

References

Atkinson, A. (1970). On the measurement of inequalitiy. *Journal of Economic Theory 2*, 244–63.

Atkinson, A. (2018). *Inequality. What can be done?* Harvard University Press.

Bosmans, K. (2016). Consistent comparisons of attainment and shortfall inequality: a critical examination. *Health Economics 25*(11), 1425–32.

Bourguignon, F. (2017). *The globalization of Inequality*. Princeton University Press.

Chakravarty, S. (2009). *Inequality, Polarization and Poverty. Advances in Distributional Analysis*. Springer.

Chakravarty, S., N. Chattopadhyay, and C. D'Ambrosio (2015). On a family of achievement and shortfall inequality indices. *Health Economics 25*(12), 1503–1513.

Clarke, P., U. Gerdtham, M. Johannesson, K. Bingefors, and L. Smith (2002). On the measurement of relative and absolute income-related health inequality. *Social Science and Medicine 55*, 1923–8.

Efthimiou (2010). *Introduction to functional equations*.

Erreygers, G. (2009). Can a single indicator measure both attainment and shortfall inequality? *Journal of Health Economics 28*, 885–93.

Erreygers, G. and T. Van Ourti (2011). Measuring socioeconomic inequality in health, health case and health financing by means of rank-dependent indices: A recipe for good practice. *Journal of Health Economics 30*, 685–94.

Kenny, C. (2004). Why are we worried about income? nearly everything that matters is converging. *World Development 33*(1), 1–19.

Lambert, P. and B. Zheng (2011). On the consistent measurement of attainment and shortfall inequality. *Journal of Health Economics 30*, 214–9.

Lasso de la Vega, C. and O. Aristondo (2012). Proposing indicators to measure achievement and shortfall inequality consistently. *Journal of Health Economics 31*, 578–583.

Marshall, A. and I. Olkin (1979). *Inequalities: Theory of Majorization and Its Applications*. New York: Academic Press.

Micklewright, J. and K. Stewart (1999). Is the well-being of children converging in the european union? *Economic Journal 109*(459), 692–714.

Milanovic, B. (2018). *Global inequality. A new approach for the age of globalization*. Harvard University Press.

Piketty, T. (2015). *The economics of Inequalitiy*. Harvard University Press.

Seth, S. and S. Alkire (2017). Did poverty reduction reach the poorest of the poor? complementary measures of poverty and inequality in the counting approach. In S. Bandyopadhyay (Ed.), *Research on Economic Inequality*, Volume 25, pp. 63–102. Emerald Publishing Limited.

Seth, S. and M. McGillivray (2018). Social choice and welfare. *Composite indices, alternative weights, and comparison robustness 51*, 657–679.

Shorrocks, A. (1980). The class of additively decomposable inequality measures. *Econometrica 48*, 613–25.

Wagstaff, A. (2005). The bounds of the concentration index when the variable of interest is binary, with an application to immunization inequality. *Health Economics 14*, 429–32.

Wagstaff, A. (2009). Correcting the concentration index: a comment. *Journal of Health Economics 28*, 516–20.

Wagstaff, A. (2011). The concentration of a binary outcome revisited. *Health Economics 20*(10), 1155–60.

Appendices

Appendix A1 Proof of Proposition 1

Let us start with an $\mathbf{x} \in \mathscr{X}_n$ for some $n \in \mathbb{N} \setminus \{1\}$ such that $\mathbf{x} \neq \hat{\mathbf{x}}$. Given that a regressive transfer increases inequality, we may always perform one or more regressive transfers until exhaustion to obtain the maximum inequality distribution $\hat{\mathbf{x}} \in \mathscr{X}_n$. Since a regressive transfer does not alter the mean of a distribution, clearly, $\mu(\mathbf{x}) = \mu(\hat{\mathbf{x}})$. Thus, an MID for any $\mathbf{x} \in \mathscr{X}_n$ always exists.

Note now that each element within $\mathbf{x}$ is bounded between 0 and U by definition and so it is not possible to perform further regressive transfers once the bounds are reached.[13] Now, there can be two types of cases: (i) $\mu(\mathbf{x}) \in \mathbb{G}$ and (ii) $\mu(\mathbf{x}) \notin \mathbb{G}$, where $\mathbb{G}_n = \{0, U/n, \ldots, (n-1)U/n, U\}$ is the set of $n+1$ equally-spaced grid points between 0 and U.

Case (i): Whenever $\mu(\mathbf{x}) \in \mathbb{G}$, then there exists a non-negative integer $n' \leq n$ such that $\mu(\mathbf{x}) = n'U/n$. Clearly, $\mu(\mathbf{x}) = n_1 \times U + (n - n_1) \times 0$. Starting with $\mathbf{x}$, it is possible to have a series of regressive transfers until a distribution with n_1 elements equalling U and $n - n_1$ elements equalling zero is reached. In this case, the maximum inequality distribution $\hat{\mathbf{x}}$ is bipolar.

Case (ii): Whenever $\mu(\mathbf{x}) \notin \mathbb{G}$, then there exist a non-negative integer $n' \leq n$ such that $n'U/n < \mu(\mathbf{x}) < (n'+1)U/n$. In this case, a series of regressive transfers are possible until n_1 elements are equal to U and $n - n_1 - 1$ elements are equal to zero. Note that it is not possible for $n_1 + 1$ elements to be equal to U because $\mu(\mathbf{x}) < (n'+1)U/n$. However, when n_1 elements are equal to U, then $\mu(\mathbf{x}) > n'U/n$ and let us denote the excess amount by $\varepsilon = n\mu(\mathbf{x}) - n'U$. It is straightforward to verify that $\varepsilon \in (0, U)$. In this case, thus, $\hat{\mathbf{x}}$ is almost-bipolar, completing our proof. ■

Appendix A2 Proof of theorem 1

Consider some $\mathbf{x} \in \mathscr{X}_n$ for some $n \in \mathbb{N} \setminus \{1\}$ and the corresponding MID $\hat{\mathbf{x}}$. We already know that

$$I(\mathbf{x}) = \begin{cases} \dfrac{f(\mathbf{x}) - f(\bar{\mathbf{x}})}{f(\hat{\mathbf{x}}) - f(\bar{\mathbf{x}})} & \text{for } \mathbf{x} \neq \bar{\mathbf{x}} \\ 0 & \text{for } \mathbf{x} = \bar{\mathbf{x}} \end{cases}, \tag{A1}$$

where $\bar{\mathbf{x}} = \mu(\mathbf{x})\mathbf{1}_n$ and $f : \mathscr{X}_n \to \mathbb{R}_{++}$ is a symmetric and strictly S-convex function with $f(\mathbf{x}^S) - f\left(\bar{\mathbf{x}}^S\right) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ for some $m, c \in \mathbb{R}_{++}$. It straightforwardly follows from its formulation in Equation A1 that I satisfies the equality principle as $I(\bar{\mathbf{x}}) = 0$ and the maximality principle as $I(\hat{\mathbf{x}}) = 1$ since $f(\hat{\mathbf{x}}) - f(\bar{\mathbf{x}}) > 0$. We now show that I satisfies the rest of the properties.

First, suppose that $\mathbf{y} \in \mathscr{X}_n$ is obtained from $\mathbf{x}$ such that $\mathbf{y} = \mathbf{x}\mathbf{P}$, where $\mathbf{P}$ is a permutation matrix. By definition, $\mu(\mathbf{x}) = \mu(\mathbf{y})$ and so $\hat{\mathbf{y}} = \hat{\mathbf{x}}$ and $f(\hat{\mathbf{y}}) = f(\hat{\mathbf{x}})$. Given that f is symmetric, then $f(\mathbf{y}) = f(\mathbf{x})$. Hence, $I(\mathbf{y}) = I(\mathbf{x})$ and I satisfies *anonymity*.

Second, suppose $\mathbf{y}' \in \mathscr{X}_n$ is obtained from $\mathbf{x}$ by a regressive transfer. Also, by definition, $\mu(\mathbf{x}) =$

[13] The proof proceeds in similar line of argument as the proof of Theorem 1 in Seth and McGillivray (2018).

$\mu(\mathbf{y}')$ and so $\hat{\mathbf{y}}' = \hat{\mathbf{x}}$ and $\bar{\mathbf{y}}' = \bar{\mathbf{x}}$. Hence $f(\hat{\mathbf{y}}') = f(\hat{\mathbf{x}})$ and $f(\bar{\mathbf{y}}') = f(\bar{\mathbf{x}})$. Given that f is strictly S-convex, then $f(\mathbf{y}') > f(\mathbf{x})$. Hence, $I(\mathbf{y}') > I(\mathbf{x})$ and I satisfies the *transfer principle*.

Finally, we show that I satisfies *strong consistency* by showing that $I(\mathbf{x}) \leq I(\mathbf{y}) \Leftrightarrow I(\mathbf{x}^S) \leq I(\mathbf{y}^S)$ for any $\mathbf{x}, \mathbf{y} \in \mathscr{X}_n$. Let $I(\mathbf{x}^S) \leq I(\mathbf{y}^S)$ and so, by Equation A1,

$$\frac{f(\mathbf{x}^S) - f(\bar{\mathbf{x}}^S)}{f(\hat{\mathbf{x}}^S) - f(\bar{\mathbf{x}}^S)} \leq \frac{f(\mathbf{y}^S) - f(\bar{\mathbf{y}}^S)}{f(\hat{\mathbf{y}}^S) - f(\bar{\mathbf{y}}^S)}. \tag{A2}$$

Then, plugging $f(\mathbf{x}^S) - f(\bar{\mathbf{x}}^S) = c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m$ in Equation A2, we obtain:

$$\frac{c[f(\mathbf{x}) - f(\bar{\mathbf{x}})]^m}{c[f(\hat{\mathbf{x}}) - f(\bar{\mathbf{x}})]^m} \leq \frac{c[f(\mathbf{y}) - f(\bar{\mathbf{y}})]^m}{c[f(\hat{\mathbf{y}}) - f(\bar{\mathbf{y}})]^m} \Leftrightarrow \frac{f(\mathbf{x}) - f(\bar{\mathbf{x}})}{f(\hat{\mathbf{x}}) - f(\bar{\mathbf{x}})} \leq \frac{f(\mathbf{y}) - f(\bar{\mathbf{y}})}{f(\hat{\mathbf{y}}) - f(\bar{\mathbf{y}})}.$$

Hence, $I(\mathbf{x}) \leq I(\mathbf{y})$. Following the same procedure it is easy to prove that the reverse direction of causality is also true, i.e. $I(\mathbf{x}) \leq I(\mathbf{y})$ implies $I(\mathbf{x}^S) \leq I(\mathbf{y}^S)$. Hence I satisfies strong consistency. ■