



UNITED NATIONS
UNIVERSITY
UNU-WIDER

WIDER Working Paper 2018/141

**What might explain today's conflicting
narratives on global inequality?**

Martin Ravallion*

November 2018

Abstract: How unequal is the world today? Is global income inequality falling, as many economists claim, or is it rising, as one often hears? This paper reviews the arguments and evidence. A number of concerns about the underlying data are identified, with biases going in both directions. Conceptual issues further cloud the picture. The claim that global inequality has been falling since 1990 can be defended for a subset of the admissible parameter values, but only a subset. Global inequality is found to be rising if one or more of the following conditions holds: (i) one attaches a high ethical weight to the poorest; (ii) one has a strong ethical aversion to high-end inequality; (iii) one takes a nationalistic perspective, emphasizing relative deprivation within countries; or (iv) one sees inequality as absolute rather than relative. Popular debates on this topic would benefit from greater clarity on the concepts used, and greater awareness of data limitations.

Keywords: global inequality, measurement, household surveys, axioms, growth

JEL classification: D31, D63, O15

Acknowledgements: For their comments the author is grateful to Conchita D'Ambrosio, Marc Fleurbaey, Ravi Kanbur, Max Kasy, Paul Segal, Finn Tarp, Milan Thomas, Dominique van de Walle, and participants at the UNU-WIDER conference 'Think Development, Think WIDER', the 'Normative Ethics and Welfare Economics' conference, University of Pennsylvania, and a seminar at the School of Foreign Service, Georgetown University.

* Department of Economics, Georgetown University, Washington, DC, USA, email: mr1185@georgetown.edu.

This study has been prepared within the UNU-WIDER project on '[Inequality in the giants](#)'.

Copyright © UNU-WIDER 2018

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9256-583-1

Typescript prepared by Joseph Laredo.

The United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Finland, Sweden, and the United Kingdom as well as earmarked contributions for specific projects from a variety of donors.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

The views expressed in this paper are those of the author(s), and do not necessarily reflect the views of the Institute or the United Nations University, nor the programme/project donors.

1 Introduction

Public attention to the global distribution of income is probably greater today than at any prior time in history.¹ Yet we hear two very different narratives about global inequality. Many economists claim that inequality has been falling in the world since around 1990.² For example, a *New York Times* article by Cowen (2014) carried the headline ‘Income Inequality Is Not Rising Globally. It’s Falling’.³ This trend has been seen to be driven by falling inequality between countries—the disparities in their mean incomes relative to the global mean. Some observers have been led to anticipate a far more equal world ahead. In this vein, Beddoes (2012) writes: ‘The gap between the world’s rich and poor will be far narrower in 2050’.

This stands in marked contrast to the claims one often hears about the rising gap between the world’s rich and poor. There are many examples. The website of Oxfam International⁴ refers to ‘A world getting more unequal’ and an ‘inequality crisis’. Similarly, in a book on global inequality, Hickel (2017: 16) writes that ‘inequality has been exploding’. Many observers point to the new ‘super-rich’. For example, as evidence of today’s huge global inequality, Basu (2018) observes that ‘the three richest persons have more wealth than all people of three nations—Angola, Burkina, Congo DR.’

Deciding which (if either) of these narratives one believes is clearly an important aspect of how one evaluates overall social progress, and the efficacy of existing economic and political institutions. Here too one hears different views. For example, Cowen (2014) argues that falling global inequality is a sign of our success, and reduces the need for public redistribution, while Hickel (2017) sees rising global inequality as indicative of a development delusion perpetuated by international financial institutions striving to justify global capitalism.

This paper tries to make some sense of these differing narratives. One might dismiss one side or the other as poorly informed, or as some ideologically driven conspiracy to hide the truth. But these are not very satisfactory responses. Both sides have their data; indeed, their sources are mostly the same. Conspiracies to delude are unlikely to work for long. The paper tries to probe more deeply into the foundations of current evidence and debates. The focus is on income inequality. Of course, there are also inequalities in wealth, and in ‘non-income’ dimensions of welfare, such as health, education, rights, and freedoms. But income inequality is the obvious place to start, given the attention it receives. The paper’s intended audience is not specialists on measurement, but economists and others in the public at large using and interpreting data on inequality.

The paper demonstrates that the view one takes of global income inequality—the stylized facts one identifies—can be highly sensitive to relaxing some of the (often implicit) assumptions made in measurement. Those assumptions relate both to how one deals with certain systematic data deficiencies, and to the concept of ‘inequality’ one uses, including its ethical premises. The

¹ Evidence for this claim can be found by entering ‘global inequality’ and ‘global poverty’ in the Google Ngram Viewer; [here](#) you will find the plot of the incidence of these phrases in all digitized text since 1950.

² See Anand and Segal (2017), Bourguignon (2016), Milanovic (2016), and Niño-Zarazúa et al. (2017).

³ Cowen was drawing on the evidence in a working paper, subsequently published in Lakner and Milanovic (2016a).

⁴ <https://www.oxfam.org/en/even-it/5-shocking-facts-about-extreme-global-inequality-and-how-even-it-davos> (accessed 15 November 2018).

concepts favoured by economists and statisticians can differ substantially from those of the population at large. While the paper does not come to a definitive conclusion as to which of the two narratives described above is closer to the truth, it is hoped that by making the assumptions explicit and comprehensible, it will make a more constructive debate possible.

The next section provides an overview of what we already know about global income inequality, based on the prevailing approach in economics, though with some new empirics for developing countries. Section 3 turns to issues pertaining to the underlying survey-based data, recognizing that perceptions of inequality may be sensitive to aspects of reality not adequately captured in standard data sources. The rest of the paper takes the data as given and turns instead to two conceptual issues. Section 4 examines the issue of what trade-offs one accepts among different levels of living, and points out that with sufficient ethical aversion to extreme inequality—in either tail of the distribution—one will conclude that global inequality is in fact rising. Section 5 turns to the distinction between absolute and relative inequality. The widespread use of relative measures rests on a ‘scale invariance axiom’ (SIA) that is routinely assumed by economists measuring global inequality but is unlikely to be widely endorsed by the public at large. Focusing instead on absolute inequality there can be little doubt that global inequality is rising. But this also points to a potential trade-off between absolute inequality and absolute poverty—an important trade-off that has received rather little attention.

2 An overview of the evidence on global income inequality

Looking back over 200 years, the best available evidence suggests that global income inequality was on a rising trend until about 1990 (Bourguignon and Morrisson 2002). This was mainly driven by much of today’s rich world taking off economically from the early 19th century. Indeed, average inequality within countries was stagnant or even falling over much of this period, most notably over the middle 50 years of the 20th century.

This is believed to have changed dramatically around the end of the 20th century. The same measures suggest that an overall pattern of falling inequality between countries emerged, alongside rising average inequality within countries. Figure 1 shows the series of global inequality measures from Bourguignon (2016). We see the fall in global inequality, markedly so in the new millennium.⁵ This has been driven by a decline in inequality between countries, which accounts for the bulk of total inequality.⁶ Average inequality within countries has edged upwards.


With falling global inequality and a rising global mean income, global measures of poverty are expected to fall when judged against poverty lines that are fixed in real terms over time and across countries. This is intuitive, although it is theoretically possible for any standard poverty measure to respond perversely to lower inequality alongside a rising mean (Datt and Ravallion 1992). Nonetheless, Chen and Ravallion (2010) show that the intuitive expectation is confirmed by the

⁵ Anand and Segal (2008) provide a compilation of estimates of global inequality over 1960–2000 that suggest an ambiguous picture, although most estimates show a decline in the 1990s. However, since their series only goes up to 2000, it misses the more marked decline in the new millennium (Figure 1).

⁶ Unlike the popular Gini index, for the Theil class of indices (one of which is used in Figure 1), the within and between components add up exactly to total inequality. (Exact aggregation for the Gini index only holds for non-overlapping distributions.)

data; indeed, falling poverty measures are found over a very wide range of poverty lines (and measures), up to (and beyond) the US official line.⁷

Figure 1: Global inequality and its between- and within-country components



Source: Author's illustration based on Bourguignon (2016: table 1).

'Globalization' has been seen as a major driver of these changes. A number of observers have argued that globalization has simultaneously decreased inequality between countries and increased it within them; see, for example, Bourguignon (2016) and Milanovic (2016). While the role of globalization in determining global inequality is not the topic of this paper, it can be noted that its causal role is not beyond dispute. Elsewhere, I have raised questions about the thesis that globalization has produced the pattern in Figure 1, drawing on the evidence from research on both growth and distributional changes (Ravallion 2018a).

There is much heterogeneity across countries and over time in the changes in the aggregate statistics for within-country inequality in Figure 1. Inequality has been rising in a majority of countries in the rich world, but not everywhere.⁸ In a compilation of national Gini indices, estimated on a reasonably consistent basis, Atkinson and Morelli (2014) find that inequality has been increasing in recent years for about two-thirds of the 25 countries studied. (Only 7 of the 25 are developing countries, and inequality has been increasing in 4 of those.) There appears to have been even more heterogeneity within the developing world. The developing countries with a trend increase in inequality over the last 20 years or so include the two most populous, China and India, which are clearly putting upward pressure on the (population-weighted) within-country component of global inequality, such as in Figure 1. However, inequality is falling in many developing countries; for example, there are clear signs that inequality has stabilized in China in recent years (Cai et al. 2018; Kanbur et al. 2017). There is also evidence of a process of inequality


⁷ For recent evidence on both absolute and relative poverty see Ravallion and Chen (2017). For the latest estimates of absolute poverty measures across multiple poverty lines see the World Bank's PovcalNet site (<http://iresearch.worldbank.org/PovcalNet/home.aspx>; accessed 15 November 2018).

⁸ Depending on the time period, one finds falling inequality in (for example) Belgium, France, Greece, Hungary, and Spain (Atkinson and Morelli 2014; Morelli et al. 2014; OECD 2011).


convergence across countries, with inequality tending to fall in high-inequality countries and rise in low-inequality countries, though the process appears to be slow (Ravallion 2003a).

Figure 2: Inequality within the developing world

(a) Developing world as a whole



(b) Within-country inequality by region



Source: Author's calculation based on PovcalNet.


To provide an overall description for the developing world, Figure 2a plots one of the Theil indices of inequality, namely the Mean Log Deviation (MLD) (also called Theil(0)), for the developing world only, and its within- and between-country components.⁹ The change in direction over the

⁹ See Theil (1967). For an overview of these and other inequality measures see Cowell (2000).

last few years appears to be mainly due to growth in China, which has surpassed the mean for the developing world. Panel (b) of Figure 2 gives a breakdown of MLD by region. Latin America and the Caribbean is the region with the highest average inequality among its countries, though this has been falling since the early 2000s. There has been a trend increase in average inequality among the countries of East Asia. Other regions have shown little trend either way.

These observations indicate that the idea of a common global force of economic integration driving up inequality everywhere can be readily dismissed. Inequality appears to fall in some developing countries when they are opened up to trade and grow in the aggregate, while inequality increases in other countries (Ravallion 2006). There are clearly many other forces in play. Indeed, during periods of economic growth we have seen falling inequality within developing countries about as often as we have seen rising inequality. This was first demonstrated by Ravallion and Chen (1997), and more recently by Ravallion (2004), using the longest available periods ('spells') between two national surveys for the same country with the same welfare indicator (either consumption or income). Using survey data up to the late 1990s, Ravallion (2004) found for 120 spells that the simple correlation coefficient between proportionate changes (annualized difference in logs) in the Gini index and those for mean household consumption or income was -0.06. Figure 3 provides an update to Ravallion (2004) using an extra 10 or more years of survey data, thus capturing the higher growth rates we have seen since 2000; the median date of the second survey in the 144 spells is 2012. Inequality increased in about half the spells (70/144 for the Gini index and 68/144 for MLD). The update in Figure 3 indicates a small positive correlation ($r=0.18$),¹⁰ which is statistically significant at the 5 per cent level.

Figure 3: Relative inequality and growth in household income per capita



Source: Author's calculations based on Ravallion (2004).

So there is more sign that the higher growth rates of mean household income seen in the new millennium have often come with increases in income inequality in developing countries. However, it is hardly a strong feature of even the new data in Figure 3. There are many instances of falling

¹⁰ This also holds if one uses changes in MLD ($r=0.20$).

inequality in growing economies. Indeed, inequality is falling in half (59/119) of the countries with positive growth in the mean in Figure 3. There is clearly a lot more to the story of what is driving the country-level distributional changes observed in the developing world.

This paper will not go any further into explaining these changes in inequality at the country level. Rather the focus will be on the numbers on global inequality, including both their statistical veracity and their conceptual basis.

3 Data concerns

There are many data issues, related to the way consumption and income are measured in practice, the design of the household surveys used, the price indices, and the census data.¹¹ This discussion focuses on those issues that are likely to be important to global inequality measures.

Let us begin by summarizing the standard practices underlying the estimates in the previous section. Almost all household surveys use personal interviews.¹² The household data refer to either consumption expenditure or disposable income, as reported by respondents for stipulated (often rather short) recall periods. Standard practice by statistics offices is to use a survey instrument that can cover all income sources and/or market goods and services consumed, including imputed values for consumption from own production, as is important for farm households. However, while cash transfers received from the government are included in the income aggregates, it appears to be rare to include imputed values for in-kind public services consumed. Consumption is used more often in developing countries, while income is more common in rich countries and Latin America. (Consumption is often preferred to current incomes, as the latter tend to be more variable over time, especially in rural-based economies. Given that such variability is to some degree predictable, consumption will be a better indicator of current welfare.) This paper often refers to ‘income’ as an abbreviation for either consumption expenditure or disposable income (though it will be noted when the distinction matters).

It is well recognized that household income is not a sufficient statistic for average economic welfare within the household. The aim is generally to obtain a monetary metric of the typical welfare level within a household. The most common method in practice is to divide household income (or consumption) by a deflator that reflects both prices and demographics (household size and composition).¹³

¹¹ Further discussion of the data issues can be found in (inter alia) Anand and Segal (2008), Gottschalk and Smeeding (1997), and Ravallion (2016a: part 2).

¹² The main option is phone interviews, though these are mainly found in rich countries. All the data issues discussed here apply to phone interviews. There is evidence (for Austria) that phone interviews impart an additional (downward) bias to inequality measures (Fessler et al. 2018).

¹³ The deflator is interpretable as a poverty line (the personal cost of a fixed reference utility level) and the ratio is variously called real income (up to a scalar) or the welfare ratio (Blackorby and Donaldson 1987). This is only an exact money metric of utility under homothetic preferences. More generally, one requires an equivalent income function giving the monetary income needed to attain the current utility level at fixed reference prices and household characteristics (King 1983). This idea can also be applied with heterogeneous (non-homothetic) preferences (see, for example, Fleurbaey and Blanchet 2013). Applications are still scarce. An example in the context of measuring poverty can be found in Ravallion and van de Walle (1991).

3.1 Surveying errors

There are sampling errors of course, though the national surveys used in measuring global inequality have large samples, and aggregating globally will probably further reduce the mean error. However, there are reasons to suspect that the within-country component in Figure 1 is systematically underestimated due to non-sampling errors.

An important source of bias is selective compliance with the randomized assignments done in sampling, such that there is a lower probability of rich people being included in the final sample used to estimate inequality measures. Such non-random compliance in surveys is a concern almost everywhere, and may be becoming more worrying over time. Traditional survey instruments are time consuming, often requiring many hours and often including multiple household members. With rising living standards, the opportunity cost of time completing surveys is likely to rise; yet surveys do not appear to be getting any shorter. Higher income households are likely to have a higher marginal cost of survey participation given the opportunity cost of time. It can also be hard to reach rich households due to gated communities and/or residences.

In theory, such selective compliance has ambiguous effects on the Lorenz curve and so does not necessarily imply that inequality measures are underestimated (Korinek et al. 2006). However, the evidence we have suggests that there is a bias. On re-weighting the data for the USA, Korinek et al. (2006) estimate that correcting for selective compliance adds 5 percentage points to the Gini index, bringing it from 0.45 to around 0.50. This method requires a common support assumption; this fails if (for example) none of the super-rich participate in the survey.

There are also concerns about the under-reporting of incomes even when there is a response, especially income from capital and illicit income sources. Under-reporting too has theoretically ambiguous effects on standard inequality measures; if someone living well above the mean under-reports their income in the survey, this will not necessarily reduce the popular Gini index (and may even increase it), though such under-reporting will always reduce the MLD.¹⁴

Reflecting both selective compliance and under-reporting, estimates using income tax records have indicated larger ‘high-end’ incomes than found in surveys (Atkinson et al. 2011). Unlike household surveys, income reporting for tax purposes is required by law, often with penalties for non-compliance or false reports. However, the narrow coverage of income tax schemes in most developing countries makes this a less convincing approach in such settings, though the data from tax records that do exist can still provide a useful clue. Drawing on income tax records, Chancel and Piketty (2017) provide estimates (requiring many assumptions) of inequality measures for India that suggest considerably higher income inequality than that underlying the numbers for South Asia in Figure 2. Similarly, the finding in Figure 2 of relatively low inequality in the Middle East and North Africa, based on household surveys, changes when one draws on these other data sources to better reflect the high end; see World Inequality Lab (2018) using its World Inequality Database (WID).¹⁵ Anand and Segal (2017) estimate global inequality measures by combining survey-based distributions with data on top incomes from the WID. Then the decline in global inequality is only evident from 2005, with little sign of a trend either way prior to that.

My expectation is that inequality within countries is both higher and rising more than the data in Figure 1 suggest, on the presumption that many newly affluent respondents are reticent to fully

¹⁴ This follows from a result in Cowell and Flachaire (2018), namely that, when comparing two distributions that differ in one person’s income, the greater the distance from equality, the higher the inequality.

¹⁵ <https://wid.world/> (accessed 15 November 2018).

reveal their gains or even to participate in surveys. It is unclear what bias if any would be imparted to Figure 3.

3.2 Two neglected sources of bias

Two data issues that are relevant to inequality measures stand out as especially neglected in the literature. The first relates to intra-household inequality. The standard (indeed, near universal) assumption is that there is equality within the household. This is almost certainly wrong, and the direction of bias is clear: we will underestimate overall inequality. Evidence on the magnitude of this bias is scarce. In one of the few cases in which we can estimate the distribution of consumption within households, Lambert et al. (2014) find a Gini index in Senegal of 0.60 when one attributes to each person the average per capita consumption of his or her household. If, instead, each individual is attributed the per capita consumption of his or her subgroup within the household, then the Gini index rises to 0.63.¹⁶ It is unclear how this source of bias would affect the time profile of inequality, though there is some evidence that gender inequalities are declining in some relevant domains, notably schooling (World Bank 2011). Then we might conjecture that the impact of this omission is becoming less important over time. If so, then this source of bias may attenuate the rise in average inequality within countries.

The second issue relates to prices. Differences in prices between countries are dealt with using purchasing power parity (PPP) rates of exchange. Since price levels tend to be higher in richer countries (especially due to higher wage rates, implying higher prices of non-traded goods), using PPPs rather than official exchange rates tends to reduce the level of inequality *between* countries. For example, Milanovic (2005) shows that the global Gini index using PPPs is 0.65, compared with 0.81 using official exchange rates. PPPs are systematically revised at times in the light of the new price surveys across countries (as used to estimate the PPPs) and methodological changes. The PPPs from the 2011 round of price surveys saw upward revisions in the estimated real incomes for some developing countries (notably in Asia), implying lower global inequality measures (Inklaar and Rao 2017). It is well known that the price surveys that feed into the PPPs are biased toward urban areas, but the extent of this bias has varied across regions of the world. Ravallion (2018c) argues that the 2011 price surveys were less urban-biased in Asia than for prior years or other regions, and that this change accounts for some share of the impact of the PPP revisions on global inequality.

It is not common to include deflators for geographic cost-of-living differences *within* countries. Within-country inequality is likely to be overestimated due to this omission. Using the data that are available on regional price differentials for Canada and the USA, Lessmann (2014) confirms that standard methods (ignoring spatial price differences) overestimate inequality measures, though the trends over time are little affected. For developing countries, however, the expectation is that spatial price differentials will be attenuated with economic development (notably through lower transport costs due to better infrastructure). Then correcting for this bias can be expected to reveal a steeper increase over time in inequality within countries.

Thus these two omissions in standard data sources—intra-household inequality and sub-national prices—point in opposite directions, in terms of both the levels and the trends over time.

¹⁶ One subgroup is the head; others comprise one or more of the wives and her children. Other evidence can be found in De Vreyer and Lambert (2018), Haddad and Kanbur (1990), and Sahn and Younger (2009).

3.3 The role of nationality

In the prevailing practice among economists, ‘countries’ only have salience as arbitrary groupings of people. There is no concept of ‘nationality’ underlying Figures 1 and 2. Yet many people seem to care more about inequality within their country of residence than globally.

There are competing views on the relevance of national borders to inequality and policy responses. Philosophers such as Singer (2010) argue that national borders, or distance, are not morally relevant to the case for helping disadvantaged people whom we can help. By this view, one *should* care about everyone, no matter where they live. This has been dubbed the ‘cosmopolitan view’.¹⁷ In the literature on inequality measurement, this view is seen as an implication of the ‘anonymity’ (or ‘symmetry’) axiom, which says that it does not matter who has which income. This is an ethical premise of standard global inequality and poverty measures. Against this view, there is a nationalistic approach whereby ‘global’ inequality is simply the average level of inequality in the world. For example, this is how Eurostat (2015) measures overall inequality in Europe. Brandolini and Carta (2016) postulate a social welfare function that treats people equally within the country of residence but puts lower weight on foreigners.

This paper takes the cosmopolitan view. This does not deny that nations exist and that their governments typically take actions to address inequality within their borders. The institutional fact of nation states and the limitations of global institutions constrain what global redistribution can be achieved in practice. These real-world constraints do not, however, diminish the moral case for a cosmopolitan perspective on ‘global inequality’—a perspective that values all people of the world equally, no matter where they may happen to live.

The cosmopolitan view still allows a role for nationality, independently of ‘own-income’. The issue can be thought about in terms of omissions/errors in measuring individual welfare. Inequality and poverty measures are summary statistics of a distribution of money-metrics of welfare. The type of global inequality measure found in this literature implicitly characterizes individual welfare in a rather narrow way, as solely a function of individual consumption or income as measured in surveys. One way that national income may matter independently of own-income stems from the longstanding idea of relative deprivation. This postulates a welfare loss from economic gains to (say) co-residents that are not shared with the person in question. Then we can rationalize a nationalistic view that ‘global inequality’ is just the average national inequality across countries. This emerges as the limiting case in which only one’s relative income within the country of residence matters.¹⁸

Against this view, it can be argued that there are also (positive) external welfare gains from living in a richer country that would not be reflected in survey-based own-incomes. A case in point is Wagner’s Law (Musgrave 1969), namely that the share of national income devoted to public services rises with income.¹⁹ While Wagner’s Law does not necessarily apply to all types of public spending or all countries, it is a plausible assumption that richer countries have better public goods and that these deliver gains in economic welfare that are not adequately captured in the survey-based measures of current disposable income or consumption expenditure used in measuring global inequality. Administrative and judicial capabilities tend to improve, creating more secure

¹⁷ See the discussions in Brandolini and Carta (2016), Caney (2005), and Nagel (2005).

¹⁸ Note that the national mean of the ratio of own-income to the national mean is unity for all countries, leaving no inequality between countries.

¹⁹ For evidence of this see Afonso and Alves (2017) and Akitoby et al. (2006).

economic opportunities that are not necessarily well reflected in current incomes. Mean income can also pick up income opportunities not reflected in the recall periods used in surveys. Studies of global subjective welfare suggest that people feel better off in richer countries at a given level of own household income; one such study concludes that ‘a richer person in a rich nation would be better off than a rich person in a poor country’ (Diener et al. 2013: 273).

In short, one can point to plausible arguments and some evidence to support the view that there are *positive* external effects of living in a richer country at given own-income. This can stem from the likely positive correlation between national income and factors conducive to a higher long-run personal income, better public services, and greater security. None of these gains is likely to be properly reflected in current incomes as measured in surveys.

The implication is clear: the (large) differences in average incomes found between rich and poor countries create an extra (horizontal) inequality between their residents, not reflected in their observed current incomes. This is a source of downward bias in prevailing measures of the between-country component of global inequality. Yet the likelihood that living in a richer country delivers gains to economic welfare that are not reflected in survey-based incomes has been entirely ignored by past measures of global inequality.

What does this imply quantitatively for measures of global inequality?

3.4 Testing sensitivity to allowing national income to matter

To test the sensitivity of global inequality measures to allowing national income to matter to individual welfare one needs to adjust survey-based incomes. It is assumed that this adjustment does not change (relative) inequality within countries. All incomes within a given country are then multiplied by a constant that depends on the national mean, preserving the distribution within countries. Specifically, the adjusted income is $m_{jt}^\alpha y_{ijt}$, where $y_{ijt} (> 0)$ denotes the surveyed income of household i in country j at time t , m_{jt} is the corresponding mean in j and α is a parameter reflecting the extra value attached to (log) national income.²⁰ Global inequality is then measured using the distribution of these adjusted incomes. The standard approach among economists (as used in the estimates reported in Section 2) has $\alpha = 0$. When $\alpha = -1$ we have the strongly relative view of Easterlin (1974) and others, whereby only relative income matters ($y_{ijt}^* = y_{ijt}/m_{jt}$). (A value of $\alpha < -1$ can be ruled out under the assumption that y_{ijt}^* is non-decreasing in own-income at a given value of relative income.) For $-1 < \alpha < 0$, the welfare metric depends positively on both own-income and relative income.²¹ Then the adjusted (log) income is a weighted mean of log survey-based income and log relative income (i.e. the log of adjusted income is $(1+\alpha)\ln y_{ijt} - \alpha \ln(y_{ijt}/m_{jt})$). However, we can also allow that (on balance) a higher national mean implies higher real income at given own-income, based on the surveys, i.e. $\alpha > 0$. Clearly this will yield a higher between-country component of global inequality. The higher the value of α the greater between-country inequality (Ravallion 2018b).

A clue to the value of α can be found in the literature on subjective welfare. In most of that literature, either the samples are micro-data drawn for one country, or the analysis is based on


²⁰ Alternatively, the effect of national income could be additive ($y_{ijt}^* \equiv y_{ijt} + \alpha m_{jt}$). There are two reasons to question this. First, it does not seem plausible that the gain from living in a rich country is constant for everyone in that country; it would seem more plausible that it is greater for higher-income households. Second, the additive form does not allow for relative deprivation, whereby adjusted income has a positive weight on y_{ijt}/m_{jt} .

²¹ This is the special case considered by Milanovic and Roemer (2016) and Ravallion and Chen (2013).

comparisons of means across countries. What we need is a study of ‘global’ micro-data. One such study is Helliwell et al. (2010), who report regression coefficients of subjective well-being on own-income and national income (GDP per capita), both in logs; the ratio of the coefficient on log national income to that on log own-income gives an estimate of α .²² The regressions suggest a positive value with an upper bound estimate around $\alpha = 0.5$. However, there are also indirect effects of national income through the other control variables used in Helliwell et al., so the true value of α is likely to be higher. For example, the indirect effect via life expectancy alone would add about 0.05 to the effect of log national income on satisfaction with life, which would raise the upper bound estimate for α to around 0.6 (Ravallion 2018b).

How much does the choice of α matter empirically? Clearly the strong form of relative deprivation theory ($\alpha = -1$) implies that global interpersonal inequality is far lower than prevailing measures suggest, since it is then entirely within countries. This changes dramatically when one attaches a positive value to national income (at given own-income), such as when living in a richer country brings benefits in terms of access to non-market goods and services, and better opportunities for private support. Figure 4 shows how the global MLD varies with α for 1993 and 2012. Global inequality falls as long as $\alpha > -0.6$.²³ Also notice that the upward adjustment to the measure of global inequality rises sharply with higher α .

Figure 4: Global inequality for various weights on (log) national-mean income



Source: Author's calculation based on Ravallion (2018b).

For even moderate (positive) α , global inequality is far higher than prevailing measures suggest, and far higher than is found in the most unequal country (Ravallion 2018b). Indeed, the differences in levels of inequality due to even rather modest differences in how one values national-mean income tend to swamp the differences seen over time in standard measures, or the differences we see between countries. They are also large relative to the impact of even a substantial underestimation of the incomes of the rich. Suppose, for example, that incomes of all the richest

²² See also Diener et al. (2010) and Helliwell (2008).

²³ The intersection point is at $\alpha = -0.66$.

1 per cent in the world are actually double the numbers in Lakner and Milanovic (2016a) for 2008.²⁴ This would add about 0.1 to the MLD,²⁵ which is about the same as adding 10 per cent of log national-mean income to log own-income to allow for the gains from living in an economically better-off country.

In summary, the stylized fact that overall inequality has been falling since around 1990 is not robust, though one only finds rising inequality with a seemingly high negative weight on national income, such as due to a strong welfare effect of relative deprivation. The finding of falling between-country inequality since 1990 is robust whatever value (positive or negative) one attaches to national income in assessing individual economic welfare.

This discussion has focused on data-related issues. The rest of the paper will largely take the data as given and focus on the conceptual foundations of prevailing measures.

4 Ethical aversion to extremes in either tail

It has long been recognized that an inequality index can be thought of as a summary statistic of the normative judgements made about how different levels of income are weighted in assessing social welfare. This was made explicit in Dalton (1920) and developed further by Atkinson (1970), who proposed a class of inequality measures in which the aversion to inequality is represented by an ethical parameter, reflecting the trade-offs allowed between incomes at different levels. (The discussion returns to this measure.)


We must first unpack Figure 1 to see how income gains were distributed by level of income. Milanovic (2013) and Lakner and Milanovic (2016a) provide an informative picture of the evolution of income distribution in the world. They plot the proportionate gain in income over 1988–2008 against fractiles of the income distribution, as reproduced in Figure 5; this is a version of a ‘growth incidence curve’ (GIC) (Ravallion and Chen 2003).²⁶ Figure 5 has been dubbed the ‘elephant chart’, since it traces the shape of an elephant’s head with its trunk held high. On the right side we see rising inequality in the rich world; between the 80th percentile (from the bottom) and the top 1 per cent globally we see a steeply positive curve (the elephant’s raised trunk), rising from near zero growth to over a 60 per cent gain for the top percentile. But we also see something equally striking—the marked proportionate rise in incomes for those near the middle of the global distribution (the elephant’s massive and expanding head). This came with considerably slower growth for the poorest decile.

²⁴ Lakner and Milanovic estimate that in 2008 the world’s richest 1 per cent had an average income of US\$64,213 (converted at PPP for 2005), while the overall mean was US\$4,097.

²⁵ Let all incomes of the richest p^r proportion of the population, with income share s^r , be underestimated by a factor k . Then the change in MLD is $(s^r - p^r) \ln k$.

²⁶ The methodology used to construct the GIC in Figure 4 is explained in Lakner and Milanovic (2016a). Note that the version of the GIC in Lakner and Milanovic gives growth rates for ventiles (with the top 1 per cent separated out) rather than percentiles. This smooths their curve. The percentile version can be found in Corlett (2016). This shows negative growth rates among the poorest and in a neighbourhood of the 80th percentile. These have been averaged out in the Lakner and Milanovic version, as also in Milanovic (2016). The negative values at the bottom probably reflect compositional effects, given that the set of countries is not held fixed. This is consistent with the fact that the quasi-non-anonymous GIC in Lakner and Milanovic (2016a: figure 5) does not show any negative growth rates.

Figure 5: The elephant graph of Lakner and Milanovic



Source: Author's illustration based on estimates in Lakner and Milanovic (2016a).


Based on the elephant graph, Milanovic (2016) argues that the rich world's lower-middle class—by interpretation, those living around the 80th percentile of Figure 5—have seen little or no gain from globalization. This is in marked contrast to the middle class of the developing world, who have seen substantial gains in the wake of the falling incidence of absolute poverty.²⁷ The largest percentage gain in the elephant graph is close to the global median. In Milanovic's interpretation, the emerging middle class in the developing world have been the big gainers from globalization, while the losers were the (relatively) poor and middle class within the rich world.

Whether or not one agrees with Milanovic on the importance of globalization to the elephant graph, it is clear that this is a much more ambiguous picture of distributional change than suggested by the claim that 'global inequality is falling'. The global Lorenz curves intersect internally, both at low percentiles and in the second decile from the top, as can be seen in Figure 6. While the overall Gini index fell (from 72 per cent to 71 per cent), this came with a marked inward shift of the Lorenz curve around the middle—between the 30th and 80th percentiles—and an outward shift among the top decile, and a declining share for the poorest 5 per cent. The Lakner-Milanovic estimates imply that the share of the world's top 1 per cent rose from 12 per cent to 15 per cent between 1988 and 2008. Some valid inequality measures (such as the Gini index and the Theil index, as in Figure 1) can show a decrease, while other equally valid measures do not. This is an implication of the fact that there is not Lorenz dominance in Figure 6.²⁸

²⁷ Ravallion (2010) argues that the developing world's middle class can be defined as those who are not poor by typical standards in poor countries, but are still poor by the standards of rich countries. The middle class has expanded greatly, and was around 50 per cent in 2005 (by the specific measure used in Ravallion 2010), and is undoubtedly higher still today.

²⁸ Lorenz dominance requires that one Lorenz curve is entirely within the other up to its end points. Such dominance implies an unambiguous ordering for all measures satisfying the Pigou-Dalton transfer axiom (Atkinson 1970).

Figure 6: Lorenz curves for global income 1988 and 2008



Source: Author's illustration based on estimates in Lakner and Milanovic (2016a).

Consider now the Atkinson (1970) index. This has a parameter ε reflecting the aversion to inequality; a higher value of ε implies that one is willing to incur a greater loss when transferring money from the rich to the poor (i.e. a lower share actually reaching the poor) and yet still judge that social welfare has increased. More precisely, the Atkinson index can be written as

$$1 - \left(\frac{1}{n} \sum_{i=1}^n \left(\frac{y_i}{\bar{y}} \right)^{1-\varepsilon} \right)^{1/(1-\varepsilon)} \quad (1)$$

where y_i is the income of person $i=1, \dots, n$ while the overall mean is \bar{y} , and where $\varepsilon > 1$ is the ethical parameter reflecting inequality-aversion;²⁹ the higher the value of ε the greater the loss one is willing to incur when transferring money from the rich to the poor and yet still judge that social welfare has increased. In Ravallion (2018a) I calculate that the Atkinson index of global inequality fell over 1988–2008 for $\varepsilon \leq 4$ but that inequality rose for $\varepsilon = 5$. This is a high value for the Atkinson parameter compared with those found in the literature, which are rarely above 2. (Atkinson's (1970) illustrative calculations of his index tested sensitivity up to 2.5.) The upshot of these observations is that with sufficiently strong aversion to inequality, one will judge that global inequality has risen over this period.

Like the Gini and Theil indices, the Atkinson index satisfies the Pigou-Dalton transfer axiom, namely that any (mean-preserving) transfer for which the donor has a higher income than the recipient reduces inequality. This axiom need not be universally accepted. In a survey I did in 2018 of my undergraduate students (using a confidential computer-based survey tool) I asked which of these two distributions had higher inequality: A: (2, 4, 6, 10) or B: (2, 5, 5, 10). The transfer axiom says A but 31 per cent of my students said B. I also asked the same question using the anonymous

²⁹ For $\varepsilon = 1$ the Atkinson index is 1 minus the ratio of the geometric mean to the arithmetic mean.

polling tool in Twitter. Out of 214 respondents in a 24-hour period, I again found that about half (53 per cent) ranked the two distributions consistently with the transfer axiom. By contrast, 21 per cent said B was more unequal, while the remainder said that the extent of inequality was no different between A and B.

In exploring this further with my students, I found that almost all (92 per cent) said that (2, 4, 6) was more unequal than (3, 4, 5), consistently with the transfer axiom.

So it is not that those who thought that (2, 5, 5, 10) was more unequal than (2, 4, 6, 10) generally rejected the transfer axiom; rather (as came out in class discussion), they put extra weight on the greater inequality in the tails, notably among the ‘rich’. The fact that the richest person, with an income of 10, had relatively more than the next richest was a troubling feature of B to my students, and outweighed the more equal middle. This is an example of what can be called an ethical aversion to ‘high-end inequality’, and it may override the Pigou-Dalton principle, and not be reflected in any standard inequality index.

Of course, the fact that growth is positive for such a large segment of the population in Figure 5 is good news. This ensures that we see a reduction in any standard measure of absolute poverty (for which the poverty line is fixed in real terms over a wide range of possible poverty lines). Indeed, one finds first-order dominance over this period, for a very wide range of poverty lines, up to and beyond the US official line (Chen and Ravallion 2010, 2013).³⁰

However, there has been much less progress for the world’s poorest, who can reasonably be said to have been ‘left behind’. We already saw a hint of this in the elephant graph (Figure 5). But this graph is highly aggregated at the bottom. We need a lens with higher magnification, and we need to fix the set of countries to avoid selection bias. To see how the poorest are doing we need to measure the floor to living standards—below which their density is zero and above which it is positive. The floor cannot be reliably measured by the lowest observed consumption or income in a survey, which is likely to be a noisy indicator. Elsewhere I have proposed that the floor should be estimated instead as the weighted mean consumption of those living below some level, with higher weight on people with lower observed consumption (Ravallion 2016b). When the weights decline linearly, the expected value of the floor is $z(1 - \frac{SPG}{PG})$, where z is the income level above which there is no chance of being the poorest person, while SPG and PG are the squared-poverty gap and poverty gap indices using z as the poverty line. This measure indicates only very modest growth in the floor, which is still barely above a survival level (Ravallion 2016b).


For the purpose of this paper, Figure 7a provides an update of the estimates in Ravallion (2016b). I have set z at the World Bank’s international poverty line of US\$1.90 per person per day at 2011 PPP. Since around 2000, the developing world has seen a substantial increase in mean consumption, but this has clearly not been shared by the poorest. We see both absolute and relative divergence between the floor and the mean. One caveat on Figure 7a is that the data include income surveys (about one third of the 2,400 household surveys on which the figure is based). Current incomes can be zero or negative, but this is probably not indicative of living standards given the scope for saving and (possibly) borrowing. However, the estimated floor rises only

³⁰ First-order dominance implies an unambiguous poverty ranking for all additive measures (Atkinson 1987). Note that first-order dominance is not implied by the Lakner-Milanovic GIC. However, this reflects the fact that the set of countries is not fixed; when one holds countries constant, one obtains positive growth rates at all percentiles in their data set (Lakner and Milanovic 2016a: figure 5).


slightly if one takes out the income surveys and only uses consumption, as can be seen in Figure 7b.

Figure 7: Consumption floor for the developing world

(a) Using both income and consumption surveys



(b) Separating out consumption surveys



Source: Author's calculation based on Ravallion (2016b).

This finding, that the poorest have seen rather little progress, raises a serious moral concern about how the distribution of income is evolving in the world. The concern echoes social policy discussions, which have often put emphasis (at least in their rhetoric) on the need to raise the floor. For example, the desire to 'leave no one behind' is prominent in the United Nations Sustainable

Development Goals (SDGs)³¹, although these results suggest that the first SDG of eliminating extreme poverty by 2030 is a long way from being attained, given that the floor would need to rise to the level of the poverty line. There is a (statistically significant) positive slope over time to how the floor has evolved, but the rate of increase is clearly very small (Figure 7). At this rate of progress, extreme poverty will not be eliminated for another 200 years or more (Ravallion 2018d). In the period from the middle of the 19th century to the middle of the 20th, during which today's rich world virtually eliminated extreme absolute poverty, more progress appears to have been made in raising the consumption floor than we are seeing in the developing world today (Ravallion 2016b).

We have seen that introducing a stronger ethical aversion to inequality changes the assessment of whether global inequality is increasing. If one puts a very high weight on the poorest, one will also conclude that global inequality is rising. The same is true if one introduces greater aversion to rising high-end inequality. These two factors—the lack of progress for the poorest, and steep gains to the richest—thus point in the same direction.

The findings of this section also point to the limitations of some widely used inequality measures. As noted, the fact that there is not Lorenz dominance implies that some measures will show a decrease in global inequality and some will show an increase. This relates to the underlying differences in the ethical weights attached to changes at different income levels. Measures such as the ratio of the mean to the median and the inter-quartile range may be quite insensitive to what is happening at the tails, among the poorest and the richest, though (by the same token) they will be more robust to errors in the tails than other measures.

5 Absolute inequality

The empirical literature has focused almost solely on measures of relative inequality, whereby the measure depends on the ratios of incomes. This class of measures follows from the scale independence axiom (SIA) in the theory of inequality measurement, which says that the measure of inequality does not change when all incomes are multiplied by a constant. However, this *is* an axiom. It need not be accepted. The alternative axiom is translation invariance, which says that the inequality index is unaffected by adding a constant. This yields absolute inequality measures that depend instead on the absolute differences (not normalized by the current mean). The choice depends entirely on which axiom one prefers—scale invariance or translation invariance. There is no right or wrong answer, as theoretical papers on inequality measurement have long recognized.³² Yet the bulk of the applied work on global inequality has used relative measures.³³ Indeed, this is typically done without even noting the fact that the option exists of using absolute measures.

Two examples illustrate the difference. Consider first the Gini index. The absolute Gini index is simply the average absolute difference between all pairs of incomes, $\sum_i \sum_j |y_i - y_j| / 2n^2$ (one can normalize by a fixed reference mean). By contrast, in calculating the relative Gini index all household incomes are normalized by the current mean. The second example is the standard

³¹ <https://www.un.org/sustainabledevelopment/sustainable-development-goals/> (accessed 15 November 2018).

³² See, for example, Blackorby and Donaldson (1980), Bosmans and Cowell (2010), Chakravarty and Tyagarupananda (1998), Dalton (1920), and Kolm (1976).

³³ Exceptions are Anand and Segal (2015), Atkinson and Brandolini (2010), Niño-Zarazúa et al. (2017), and Ravallion (2003b, 2004, 2014).

deviation, $s = (\sum_i (y_i - \bar{y})^2)^{1/2}$. This is clearly an absolute inequality measure, since adding a constant to all incomes leaves it unchanged. (The same is true of the variance, s^2 .) Consider instead the coefficient of variation, s/\bar{y} . This is the corresponding relative inequality measure, in that multiplying by a constant leaves the measure unchanged.


There is no sign of a popular consensus on this issue. A number of experiments (all with university students to my knowledge) have found that 40–60 per cent of participants (in the UK, Israel, Germany, and the USA) think about inequality in absolute rather than relative terms.³⁴ In the aforementioned surveys of my students at Georgetown I found that the majority do not accept the SIA. For example, when asked which of the two income distributions, (1, 2, 3) and (2, 4, 6), has higher inequality (if either), 258 (56 per cent) out of 460 students said it was the latter distribution. When comparing (2, 4, 6) with (4, 8, 12), 57 per cent said that the latter had higher inequality. Similarly, in my aforementioned Twitter survey, out of 247 responses in a 24-hour period (21/22 November 2018), 48 per cent said that (2, 4, 6) had higher inequality than (1, 2, 3), while 45 per cent said inequality was the same. While these are hardly random samples drawn from any well-defined population, they are at least consistent with the view that many people view inequality as absolute not relative.

Even if relative inequality does not change during a period of growth in mean income, the absolute income gains to the rich will obviously be greater than those to the poor, given existing inequality. We saw in the previous section that a focus on the poorest suggests rising inequality, both relative and absolute. Figure 8 gives the absolute GIC—the US\$ gains across percentiles of the whole distribution. Compared with Figure 5, the elephant’s head has shrunk greatly relative to the trunk. Over this 20-year period, the absolute gain in mean daily income of the poorest 5 per cent was 7 cents per person, while for the richest 1 per cent it was almost US\$70 (and the latter number could well be an underestimate, as noted in Section 2). In absolute terms, the developing world’s middle class and (especially) its poor have gained rather little; it is only because they started off so poor that the elephant’s head is so large in the (relative) GIC in Figure 5. While the relative Gini index fell slightly (from 0.72 to 0.71), the absolute version rose appreciably (from 0.72 to 0.90).³⁵

³⁴ The literature on survey-based perceptions of inequality has followed Amiel and Cowell (1992), who found that 40 per cent of the university students they surveyed (in the UK and Israel) thought about inequality in absolute rather than relative terms. Harrison and Seidl (1994) report similar findings for a large sample of German university students.

³⁵ For the purpose of this calculation I have normalized the indices such that the two are equal in 1988.


Figure 8: Elephant or serpent? Incidence of the absolute income gains corresponding to Figure 5



Source: Author's calculation based on estimates made by Lakner and Milanovic (2016a).

We saw in Figure 3 that there is weak correlation between changes in the relative Gini index and growth rates. The correlation becomes much stronger when one switches to the absolute Gini index, as can be seen in Figure 9. Higher rates of growth in the overall mean come with steeper increases in absolute inequality. This is hardly surprising, given that the absolute Gini index is (up to a scalar) the relative index times the mean. However, the comparison of Figures 3 and 9 underlines the sensitivity of statements about inequality and growth to the axiomatic foundations. Switch out one axiom and the picture changes a lot. The (many) people who view inequality as absolute rather than relative will probably see sharply rising inequality in growing economies. They are not wrong; they simply have a different concept of what 'inequality' means.

Figure 9: Absolute inequality and growth in household income per capita




Source: Author's calculation based on Ravallion (2004).

When combined with the tendency for measures of absolute poverty to fall with growth in mean income, the (strong) tendency for absolute inequality to rise with that growth points to a trade-off between reducing absolute inequality and reducing poverty, as discussed in Ravallion (2005). Figure 10 plots the annualized growth rate of the headcount index of absolute poverty (using US\$1.90 a day) against the corresponding growth rate in the absolute Gini index (as used in Figure 9).³⁶ This pattern suggests that those who see inequality as absolute and give high priority to reducing it may well find themselves living in an absolutely poorer world given current growth processes. Greater clarity is needed on what trade-offs one is willing to accept between reducing absolute inequality and reducing absolute poverty.

Will we see absolute inequality start to decline at some point in the future? That depends on what absolute gaps we focus on and whether relative inequality is declining. If we are talking about the world's richest 1 per cent (say) and the poorest 1 per cent, it is plain that we will not see a declining absolute gap in the foreseeable future if recent trends continue, since we are not even seeing falling relative inequality between the two groups of people.

³⁶ The correlation coefficient of -0.37 (n=136) is significant at the 1 per cent level. If one drops the outlier with the highest rate of increase in the poverty rate (Figure 10), the correlation coefficient drops to -0.28 (n=135) but it is still significant at the 1 per cent level.

Figure 10: Plot of changes in absolute poverty against changes in absolute inequality across developing countries



Source: Author's calculation based on Ravallion (2004).

There is a range of 'middle' incomes for which recent trends do suggest declining absolute inequality over the next few decades. Compare the world's middle incomes—the 50th–60th percentiles, say (just above the global median)—with the income of the 80th–90th percentiles, i.e. the group that Milanovic (2016) identifies as the rich world's 'middle class'. The elephant graph (Figure 5) shows that the former group has seen its incomes growing strongly, at 3.6 per cent per annum, over 1988–2008, while the latter group has seen little growth (0.23 per cent per annum). The ratio of mean incomes in 2008 was 5.7 (Lakner and Milanovic 2016a: table 3). Then it can be readily verified that absolute inequality between the two will decline, though it will take 53 years for the two income levels to converge if recent trends continue.³⁷ One might question whether a near zero growth rate of the rich world's middle class is politically sustainable. Suppose instead that this income group sees a 1 per cent per annum growth rate, with all else unchanged. Then the absolute gap between the rich world's middle class and the world's middle will *rise* for the next 20 years, and only then start to fall, vanishing after about 70 years.

Such calculations should be taken with a grain of salt. They only serve to illustrate that absolute inequality is likely to persist for some time even with falling relative inequality. Indeed, with current trends, the gap will rise between the world's richest and poorest, as may the gap between the rich world's middle class and the new middle class of the developing world.

³⁷ Let t^* be the number of years required for the two incomes to equalize. Then $t^* = \frac{\ln(y_1/y_0)}{\ln(g_0/g_1)}$ where y_i is the base year income of group $i=0,1$ (where $y_1 > y_0$) and g_i is the corresponding growth factor ($1 + \text{growth rate}$).

6 Conclusions

It is important to know how robust prevailing measures of global income inequality are to relaxing the (often implicit) assumptions made in measurement. This is not just an academic question. The measures used to inform public debates should accord with popular perceptions of what ‘inequality’ means. That is almost certainly not the case at present. And this provides a clue to understanding the differing narratives one hears on what is happening to inequality.

Within-country measures based on standard data sources may well underestimate inequality, though some data deficiencies point in the opposite direction. There are many issues, but one stands out in my view: selective compliance with the randomized assignments in surveys. This is a serious and (possibly) increasing concern for measuring inequality almost everywhere. This does not necessarily imply that inequality measures are underestimated, but the evidence so far suggests that they are. Appropriate re-weighting methods can address this problem in a way that is internally consistent with the survey, but it may still miss the extremes. Triangulation with other data sources (such as income tax records when coverage and compliance are adequate) can also help correct the problem. Going forward, the technology of surveying may need to adapt if it is to produce distributional data that represent well the relevant populations. Better linkage across survey and administrative databases will also help.

When one considers the conceptual foundations of prevailing approaches, one finds further reasons to question the robustness of the claim that global income inequality has been falling over the last 30 years. With sufficient ethical aversion to lack of progress by the poorest, or to steeply rising top incomes, one concludes that global inequality has been rising. A focus on the middle of the global distribution suggests that relative inequality is falling, though it will clearly be a long time (50 or more years) before the developing world’s emerging middle class catches up on average with the rich world’s middle class.

Rising global inequality is also indicated if one holds a very strong concern about relative deprivation (or, equivalently, a highly nationalistic perspective on global inequality). However, a more serious omission in prevailing measures may well be that they do not allow for the benefits of living in a richer country at any given level of own-income as measured in surveys. With even a seemingly modest positive effect of national income on individual welfare, global inequality is far higher than we think, though still falling over time.

A credible argument suggesting that global inequality may in fact be rising is found in concerns about the commonly assumed (but contentious) scale invariance axiom (SIA). It appears that many people do not accept the relativism implied by this axiom; instead, they look to the absolute gaps between ‘rich’ and ‘poor’ in assessing inequality. Then we see clear signs of rising absolute inequality in growing economies. We also see that the world’s poorest are being left behind in the wave of higher rates of economic growth in the new millennium. While there are fewer people living near the world’s floor of living standards, that floor has risen little despite overall economic growth. A measure based on the income gaps across the whole distribution indicates that absolute inequality has been rising globally with economic growth, and that can be expected to continue for some time given the current level of inequality.

The measurement issues reviewed here are salient to the debates on globalization. Different sides in that debate often appear to hold different ideas about what ‘inequality’ means (though at times one also hears claims that have no imaginable basis in reason or fact). Those who talk about the widening gap between rich and poor appear to have in mind absolute inequality, not relative inequality. Yet one cannot say that one of these concepts is right and the other wrong; the

difference is based solely on the choice between two rival axioms in the theory of inequality measurement. The standard definition in terms of relativities can be questioned; if one does not accept the SIA, one can justifiably reject relative measures in favour of absolute ones (satisfying the translation invariance axiom). In this respect, the measurement tools used in this literature appear to be woefully inadequate for informing the public discourse about ‘inequality’.

In conclusion, the claim that global income inequality has been falling over the last few decades holds over a subset of the defensible measures, but only a subset. This ambiguity belies the (confident) claims one often hears, with one side predicting a far more equal world ahead, and the other claiming that development has failed. Given the scope for sensible people to disagree on the desirable properties of an inequality measure—and there is no scientific justification for the near monopoly of relative measures in applied work by economists—more productive debates on globalization and development might be possible if both sides better understood what concepts of inequality they are using. To talk about ‘inequality’ without making explicit whether one means absolute or relative inequality is especially problematic. The non-robustness found in this paper also points to the limitations of both a single concept of inequality and any single overall measure, such as the popular Gini index. A single measure is unlikely to be acceptable to everyone. Thankfully, the more flexible tools that are now available for representing distributional change, such as (absolute and relative) growth incidence curves, can be used to better inform public debates on this topic.

References

- Afonso, António, and José Alves (2017). 'Reconsidering Wagner's Law: Evidence from the Functions of the Government'. *Applied Economics Letters*, 24(5): 346–50.
- Akitoby, Bernardin, Benedict Clements, Sanjeev Gupta, and Gabriela Inchauste (2006). 'Public Spending, Voracity, and Wagner's Law in Developing Countries'. *European Journal of Political Economy*, 22(4): 908–24.
- Amiel, Yoram, and Frank Cowell (1992). 'Measurement of Income Inequality: Experimental Test by Questionnaire', *Journal of Public Economics*, 47: 3–26.
- Anand, Sudhir, and Paul Segal (2008). 'What Do We Know about Global Income Inequality?' *Journal of Economic Literature*, 46(1): 57–94.
- Anand, Sudhir, and Paul Segal (2015). 'The Global Distribution of Income'. In Anthony Atkinson and Francois Bourguignon (eds), *Handbook of Income Distribution, Volume 2*, Amsterdam: North-Holland.
- Anand, Sudhir, and Paul Segal (2017). 'Who Are the Global Top 1 Per Cent?' *World Development*, 95: 111–26.
- Atkinson, Anthony (1970). 'On the Measurement of Inequality'. *Journal of Economic Theory*, 2: 244–63.
- Atkinson, Anthony (1987). 'On the Measurement of Poverty'. *Econometrica*, 55: 749–64.
- Atkinson, Anthony, and Andrea Brandolini (2010). 'On Analyzing the World Distribution of Income'. *World Bank Economic Review*, 24(1): 1–37.
- Atkinson, Anthony, and Salvatore Morelli (2014). 'Chartbook of Economic Inequality'. ECINEQ Working Paper 324.
- Atkinson, Anthony, Thomas Piketty, and Emmanuel Saez (2011). 'Top Incomes in the Long Run of History'. *Journal of Economic Literature*, 49(1): 3–71.
- Basu, Kaushik (2018). 'Today's Huge Global Inequality'. Twitter, 13 October. Available at: <https://twitter.com/kaushikbasu/status/1051118562151149568> (accessed 15 November 2018).
- Beddoes, Zanny Minton (2012). The Great Levelling. CNBC (16 March). Available at: <https://www.cnbc.com/id/46683856> (accessed 12 November 2018).
- Blackorby, Charles, and David Donaldson (1980). 'A Theoretical Treatment of Indices of Absolute Inequality'. *International Economic Review*, 21(1): 107–36.
- Blackorby, Charles, and David Donaldson (1987). 'Welfare Ratios and Distributionally Sensitive Cost–Benefit Analysis'. *Journal of Public Economics*, 34: 265–90.
- Bosmans, Kristof, and Frank Cowell (2010). 'The Class of Absolute Decomposable Inequality Measures'. *Economics Letters*, 109: 154–56.
- Bourguignon, François (2016). *The Globalization of Inequality*. Princeton: Princeton University Press.
- Bourguignon, François, and Christian Morrisson (2002). 'Inequality among World Citizens: 1820–1992'. *American Economic Review*, 92(4): 727–44.
- Brandolini, Andrea, and Francesca Carta (2016). 'Some Reflections on the Social Welfare Bases of the Measurement of Global Income Inequality'. *Journal of Globalization and Development*, 7(1): 1–15.

- Cai Meng, Bjorn Gustafsson, Qinghai Li, Li Shi, Chuliang Luo, and Terry Sicular (2018). 'Inequality in the Developing Giants: China'. Paper presented at the UNU-WIDER conference 'Think Development, Think WIDER'. Helsinki.
- Caney, Simon (2005). *Justice Beyond Borders. A Global Political Theory*. Oxford: Oxford University Press.
- Chakravarty, S.R., and S. Tyagarupananda (1998). 'The Subgroup Decomposable Absolute Indices of Inequality'. In: S.R. Chakravarty, D. Coondoo, and R. Mukherjee (eds), *Quantitative Economics: Theory and Practice, Essays in Honor of Professor N. Bhattacharya*. New Delhi: Allied Publishers, pp. 247–57.
- Chancel, Lucas, and Thomas Piketty (2017). 'Indian Income Inequality (1922–2014): From British Raj to Billionaire Raj?' Working Paper 2017/11. Paris School of Economics: World Wealth and Income Database.
- Chen, Shaohua, and Martin Ravallion (2010). 'The Developing World is Poorer Than We Thought, but No Less Successful in the Fight against Poverty'. *Quarterly Journal of Economics*, 125(4): 1577–625.
- Chen, Shaohua, and Martin Ravallion (2013). 'More Relatively Poor People in a Less Absolutely Poor World'. *Review of Income and Wealth*, 59(1): 1–28.
- Corlett, Adam (2016). 'Examining an Elephant: Globalization and the Lower Middle Class of the Rich World'. London: Resolution Foundation Report.
- Cowell, Frank (2000). 'Measurement of Inequality'. In A.B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*. Amsterdam: North-Holland.
- Cowell, Frank, and Emmanuel Flachaire (2018). 'Inequality Measures and the Rich: Why Inequality Increased More Than We Thought'. Mimeo. London School of Economics.
- Cowen, Tyler (2014). 'Income Inequality Is Not Rising Globally. It's Falling'. *New York Times*, 19 July.
- Dalton, Hugh (1920). 'The Measurement of the Inequality of Incomes'. *Economic Journal*, 30(9): 348–61.
- Datt, Gaurav, and Martin Ravallion (1992). 'Growth and Redistribution Components of Changes in Poverty Measures: A Decomposition with Applications to Brazil and India in the 1980s'. *Journal of Development Economics*, 38: 275–95.
- De Vreyer, Philippe, and Sylvie Lambert (2018). 'Intra-Household Inequalities and Poverty in Senegal'. Mimeo. Paris School of Economics.
- Diener, Ed, W. Ng, J. Harter, and R. Arora (2010). 'Wealth and Happiness across the World: Material Prosperity Predicts Life Evaluation, Whereas Psychosocial Prosperity Predicts Positive Feeling'. *Journal of Personality and Social Psychology*, 99(1): 52–61.
- Diener, Ed, Louis Tay, and Shigehiro Oishi (2013). 'Rising Income and the Subjective Well-being of Nations'. *Journal of Personality and Social Psychology*, 104(2): 267–76.
- Easterlin, Richard A. (1974). 'Does Economic Growth Improve the Human Lot? Some Empirical Evidence'. In P.A. David and W.R. Melvin (eds), *Nations and Households in Economic Growth*. Palo Alto: Stanford University Press.
- Eurostat (2015). *Income and Living Conditions*. Luxembourg: Eurostat. Available at: https://ec.europa.eu/eurostat/cache/metadata/en/ilc_esms.htm (accessed 15 November 2018).

- Fessler, Pirmin, Maximilian Kasy, and Peter Lindner (2018). ‘Survey Mode Effects on Measured Income Inequality’. *Journal of Economic Inequality*, 16(4): 487–505.
- Fleurbaey, Marc, and Didier Blanchet (2013). *Beyond GDP Measuring Welfare and Assessing Sustainability*. Oxford: Oxford University Press.
- Gottschalk, Peter, and Timothy Smeeding (1997). ‘Cross-National Income Inequality: Cross-National Comparisons of Earnings and Income Inequality’. *Journal of Economic Literature*, 35: 633–87.
- Haddad, Lawrence, and Ravi Kanbur (1990). ‘How Serious is the Neglect of Intra-Household Inequality?’ *Economic Journal*, 100: 866–81.
- Harrison, Elizabeth, and Christian Seidl (1994). ‘Perceptual Inequality and Preferential Judgment: An Empirical Examination of Distributional Axioms’. *Public Choice*, 79: 61–81.
- Helliwell, John (2008). ‘Life Satisfaction and Quality of Development’. NBER Working Paper 14507.
- Helliwell, John, Chris Barrington-Leigh, Anthony Harris, and Haifang Huang (2010). ‘International Evidence on the Social Context of Well-Being’. In Ed Diener, John Helliwell, and Daniel Kahneman (eds), *International Differences in Well-Being*. Oxford: Oxford University Press.
- Hickel, Jason (2017). *The Divide: A Brief Guide to Global Inequality and Its Solutions*. London: William Heinemann.
- Inklaar, R., and P. Rao (2017). ‘Cross-Country Income Levels over Time: Did the Developing World Suddenly Become much Richer?’ *American Economic Journal: Macroeconomics*, 9(1): 265–90.
- Kanbur, Ravi, Yue Wang, and Xiaobo Zhang (2017). ‘The Great Chinese Inequality Turnaround’. Centre for Economic Policy Research Discussion Paper 11892. London.
- King, Mervyn A. (1983). ‘Welfare Analysis of Tax Reforms Using Household Level Data’. *Journal of Public Economics*, 21: 183–214.
- Kolm, Serge-Christophe (1976). ‘Unequal Inequalities I’. *Journal of Economic Theory*, 12(3): 416–42.
- Korinek, Anton, Johan Mistiaen, and Martin Ravallion (2006). ‘Survey Nonresponse and the Distribution of Income’. *Journal of Economic Inequality*, 4(2): 33–55.
- Lakner, Christoph, and Branko Milanovic (2016a). ‘Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession’. *World Bank Economic Review*, 30(2): 203–32.
- Lakner, Christoph, and Branko Milanovic (2016b). ‘Response to Adam Corlett’s “Examining an Elephant: Globalisation and the Lower Middle Class of the Rich World”’. Mimeo. Graduate Center, City University of New York.
- Lambert, Sylvie, Martin Ravallion, and Dominique van de Walle (2014). ‘Intergenerational Mobility and Interpersonal Inequality in an African Economy’. *Journal of Development Economics*, 110: 327–44.
- Lessmann, Christian (2014). ‘Spatial Inequality and Development—Is there an Inverted-U Relationship?’ *Journal of Economic Development*, 106: 35–51.
- Milanovic, Branko (2005). *Worlds Apart: Measuring International and Global Inequality*, Princeton: Princeton University Press.
- Milanovic, Branko (2013). ‘Global Income Inequality in Numbers: in History and Now’. *Global Policy*, 4(2): 198–208.

- Milanovic, Branko (2016). *Global Inequality: A New Approach for the Age of Globalization*. Cambridge, MA: Harvard University Press.
- Milanovic, Branko, and John E. Roemer (2016). 'Interaction of Global and National Income Inequalities'. *Journal of Globalization and Development*, 7(1): 109–15.
- Morelli, Salvatore, Timothy Smeeding, and Jeffrey Thompson (2014). 'Post-1970 Trends in Within-Country Inequality and Poverty'. In Anthony B. Atkinson and François Bourguignon (eds), *Handbook of Income Distribution, Volume 2*. Amsterdam: Elsevier Science.
- Musgrave, Richard (1969). *Fiscal Systems*. New Haven: Yale University Press.
- Nagel, Thomas (2005). 'The Problem of Global Justice'. *Philosophy and Public Affairs*, 33(2): 113–47.
- Niño-Zarazúa, Miguel, Laurence Roopez, and Finn Tarp (2017). 'Global Inequality: Global Inequality: Relatively Lower, Absolutely Higher'. *Review of Income and Wealth*, 63(4): 661–84.
- OECD (2011). *Divided We Stand: Why Inequality Keeps Rising*. Paris: OECD
- Ravallion, Martin (2003a). 'Inequality Convergence'. *Economics Letters*, 80: 351–56.
- Ravallion, Martin (2003b). 'The Debate on Globalization, Poverty and Inequality: Why Measurement Matters'. *International Affairs*, 79(4): 739–54.
- Ravallion, Martin (2004). 'Competing Concepts of Inequality in the Globalization Debate'. Brookings Trade Forum 2004, edited by Susan Collins and Carol Graham. Washington, DC: Brookings Institution, pp.1–38.
- Ravallion, Martin (2005). 'A Poverty–Inequality Trade-Off?' *Journal of Economic Inequality*, 3(2): 169–82.
- Ravallion, Martin (2006). 'Looking Beyond Averages in the Trade and Poverty Debate' *World Development* (special issue on 'The Impact of Globalization on the World's Poor, edited by Machiko Nissanke and Erik Thorbecke), 34(8): 1374–92.
- Ravallion, Martin (2014). 'Income Inequality in the Developing World'. *Science*, 344: 851–55.
- Ravallion, Martin (2016a). *The Economics of Poverty: History, Measurement and Policy*. New York: Oxford University Press.
- Ravallion, Martin (2016b). 'Are the World's Poorest Being Left Behind?' *Journal of Economic Growth*, 21(2): 139–64.
- Ravallion, Martin (2018a). 'Inequality and Globalization: A Review Essay'. *Journal of Economic Literature*, 56(2): 1–23.
- Ravallion, Martin (2018b). 'Global Inequality when Unequal Countries Create Unequal People'. *European Economic Review*, forthcoming.
- Ravallion, Martin (2018c). 'An Exploration of the Changes in the International Comparison Program's New Global Economic Landscape'. *World Development*, 105: 201–16.
- Ravallion, Martin (2018d). 'SDG1.1: Are We on Track to Eradicating Extreme Poverty by 2030?' Blog post, Economics&poverty.
- Ravallion, Martin, and Shaohua Chen (1997). 'What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty?' *World Bank Economic Review*, 11(2): 357–82.
- Ravallion, Martin, and Shaohua Chen (2003). 'Measuring Pro-Poor Growth'. *Economics Letters*, 78(1): 93–99.
- Ravallion, Martin, and Shaohua Chen (2013). 'A Proposal for Truly Global Poverty Measures'. *Global Policy*, 4(3): 258–65.

- Ravallion, Martin, and Shaohua Chen (2017). 'Welfare-Consistent Global Poverty Measures When Relative Incomes Matter'. NBER Working Paper 23739.
- Ravallion, Martin, and Dominique van de Walle (1991). 'The Impact on Poverty of Food Pricing Reforms: A Welfare Analysis for Indonesia'. *Journal of Policy Modeling*, 13: 281–99.
- Sahn, David, and Stephen Younger (2009). 'Measuring Intra-Household Health Inequality: Explorations Using the Body Mass Index'. *Health Economics*, 18: S13–S36.
- Singer, Peter (2010). *The Life You Can Save: How to Do Your Part to End World Poverty*. New York: Random House.
- World Bank (2011). *World Development Report: Gender Equality and Development*. Washington, DC: World Bank.
- World Inequality Lab (2018). *World Inequality Report*. Paris School of Economics. Available at: <https://wir2018.wid.world/executive-summary.html> (accessed 15 November 2018).