

LIS

Working Paper Series

No. 690

The Inequality We Want: How Much Is Too Much?

Alice Krozer

March 2017



CROSS-NATIONAL
DATA CENTER
in Luxembourg

Luxembourg Income Study (LIS), asbl

The Inequality We Want: How Much Is Too Much?

Alice Krozer¹

Abstract

A key aspect defining the contemporary income distribution is the (increasing) share the top holds compared to the rest. This paper shows that income concentration increases towards the very top of the distribution, while the shares the middle- and upper-middle-income groups hold remain stable across countries and over time. Traditional indicators less sensitive to changes at the extremes of the distribution might obscure inequality's actual dimension, and thereby help perpetuate it. To avoid this, the present paper discusses a complementary indicator for the measurement of inequality: the ratio of the income share of the top 5 per cent over that of the bottom 40 per cent. The indicator is conceptually related to the recently suggested 'Palma ratio' dividing the income share of the top 10 per cent income earners by that of the bottom 40 per cent's.

Keywords: Inequality; Income Distribution; Poverty; Palma; Gini Coefficient

JEL Classification: D31, D63, I3

¹ Alice Krozer, University of Cambridge; email: ak857@cam.ac.uk

Acknowledgements: I would like to thank Branko Milanovic for sharing his WYD 2008 dataset and Marco Lilla for assistance with the LIS microdata. I am grateful for fruitful discussions with Gabriel Palma and valuable editorial comments by Lesley Ellen.

1. Introduction

A casual look at world income distribution reveals large and persistent inequalities, both between rich and poor countries, and within countries. The severity of the problem is illustrated with a few numbers, For example, that the top 1 per cent of income earners owns roughly 40 per cent of global income, and the 85 richest individuals in the world own as much wealth as the bottom 3 billion (Oxfam, 2014). In order to make more specific statements about inequality levels within certain countries, however, (such as whether they are rising or falling, how they compare to each other, etc.) we have to measure income differences and where they lie along the distribution according to some indicator.

Conventionally, this indicator would be the Gini coefficient. However, one of the crucial questions the Gini has difficulty answering is what happens at the extremes of the distribution: how does the income of the poor develop and what happens to the share the rich hold, especially the very rich? Most of the general public on the other hand is interested exactly in this aspect of inequality; rather than weighting changes to the income distribution equally, whether they happen at the top, at the bottom or in the middle of the income scale (Hasenheit, 2014), a majority of people intuitively oppose extreme concentrations (Oxfam, 2014). In fact, most people are willing to tolerate certain discrepancies in incomes (and find these justified), as long as the differences between the top and bottom ends of the distribution are held in check. Because of the way it is constructed though, the Gini overemphasises changes in the middle of the distribution, and cannot account for changes that happen in the extremes (Cobham and Sumner, 2013a).

To counter this shortcoming, Gabriel Palma recently suggested (2011) a different indicator to measure inequality particularly in the tails of the distribution. What Cobham and Sumner (2013b) came to term 'the Palma', is the ratio of the share of total income held by the top 10 per cent of income earners over that of the bottom 40 per cent, concentrating thus on

the tails of the distribution while ‘ignoring’ the middle. Concretely, the measure states that the top 10 per cent of income earners own a share x times the size of that of the bottom 40 per cent of the population. Besides the simplicity of calculation, the additional advantage of such an indicator is this intuitiveness in its interpretation, making it a persuasive tool for policy-making.

However, it would be difficult to accept such an indicator comprising nominally only half of the population as a measure of inequality over the entire income distribution—if it were not for the fact that the actual income distribution supports this ‘exclusion’ of the middle- and upper-middle-income groups. When looking at countries’ income patterns according to deciles of the population, Palma (2006) found that across the world, the middle-income groups consistently held about 50 per cent of a country’s total income independently of the country’s developmental status, income level, political regime, location, or other particular condition. Since at the same time inequality levels across countries are vastly different (ranging from a Gini of 24.6 in Slovenia to 65.4 in South Africa), this means that these differences are determined almost entirely in the extremes of the distribution, i.e. the degree of polarisation of incomes between the top and the bottom ends of the distribution differs markedly across countries (see Palma, 2006; 2011; 2014).² Palma named these diverging dynamics the centrifugal and centripetal forces of the income distribution, creating heterogeneity in the tails while homogenising the middle. For measuring inequality, it is thus not only important to focus on the extremes because this is what most people care about, but

² Palma dubs this phenomenon the 50–50 rule, where 50 per cent of the population holds about 50 per cent of total income. At the same time, the diversity in inequality levels is explained by the way the other half of the population, i.e. the top decile and the bottom four deciles, divides the remaining 50 per cent of total income amongst its constituents. The vast differences can be observed between the Latin American and Southern African countries on the one hand, where the top income groups make multiples of the share the low income population secures -- over six times again in the case of Honduras and up to nine times in South Africa – and the Nordic and Eastern European countries on the other hand, where the elites’ share is actually smaller than that of the poor groups (between 70–80 per cent of the latter’s in the cases of Slovenia or Sweden for example).

even more so because this is what characterises the actual distribution, with the opposing forces conditioning ‘its’ behaviour.

The intuition of this concentration of income in the top decile compared to the rest of the distribution has one wondering almost inevitably as to what would happen to an indicator if this upper fraction was sliced up even more, reduced to but a few percentiles—to the ‘real’ elite of the proverbial 1 per cent. Considering the simultaneous stability of their income share in the middle groups, inspiring Palma to propose his eponymous indicator, an intriguing question becomes: how are these dynamics affected when the top is reduced to 5 per cent (1 per cent)? How is world income distribution developing with regards to Palma’s hypothesis of the ‘diverging ends’ and ‘converging middles’ where smaller shares at the top are distinguished, i.e. is the middle consolidating its stable share while the rich and poor drift further apart?

This paper aims to address these questions, on the premise that concern over inequality would be greater if the actual dimension of concentration at the top was made explicit (via an adequate indicator). Until recently, it has been impossible to ‘look inside’ the high end of the distribution, for lack of available data. The creation of new datasets now allows us to take a fresh look at this issue. In this vein, the present paper aims to: (i) analyse a specially designed global dataset dissecting the income distribution into ventiles and looking for a potential homogenous middle and heterogeneous tails in a cross-sectional manner; (ii) peek into the distribution according to the top 1 per cent via a smaller subset for which data is available; and (iii) tentatively test the hypothesis of a trend towards divergence of the tails and stabilisation in the middle over time. It will introduce an additional ratio for measuring inequality levels to supplement the Palma indicator, the 5/40 ratio, and the 1/40 ratio, and argue for why especially the 5/40 version seems a useful add-on to the portfolio of existing inequality indicators.

The paper is divided into five sections. Section 2 will discuss why we should be concerned about inequality, and particularly about the current high levels of inequality.³ Section 3 will briefly outline the methodology and describe the datasets used. Section 4 will scrutinise the results obtained from the empirical data and present some answers to the question posed above. The final section will summarise the findings and discuss the implications for policy-making, as well as sketch out some lines for future research.

Before proceeding to the contextualization of inequality, it is important to bear in mind that income is only one dimension of economic inequality, which itself is only one aspect of inequality alongside with e.g. social, political, and personal inequality dimensions. A focus on that aspect necessarily relegates other important determinants – like wealth, with a significantly less equal distribution – to lower ranks. However, labour income is still the most important source of income for the vast majority of people. Even among those in the top 10 per cent of the distribution, earnings from work account for shares of 70–85 per cent (OECD, 2014), making it one of the most important components of economic inequality. Moreover, asset inequality and income inequality are closely associated across countries (Goñi et al., 2008). Finally, data on wealth accumulation in developing countries is very patchy, making large-scale analyses on wealth distribution difficult at this point.

The second caveat of this paper is its focus on a comparison of the Palma to the Gini coefficient only, rather than including other inequality measures such as the Theil, Atkinson or Concentration Indices. This is both for reasons of scope and because the Gini is the most

³ The paper largely omits the important discussion about the determinants of inequality for reasons of scope. The long list of causal factors and the resulting, at times conflicting theories, as well as inequality's ubiquitous appearance, prompted Atkinson (1983, p.206) to state that '[t]he determination of the share of wages and profits in national income is a difficult problem. [...] it involves the whole economic system, and typically everything depends on everything else.' Some of the most prominent factors mentioned to influence inequality levels include, by themselves or in combination, growth, trade liberalisation, globalisation, public social spending, political regime-type, minimum wage, conditional cash transfer programmes, diminished skill-premium, taxation, education, broader labour market policies (job creation), history and path-dependence, remittances, or the emergence of specific sectors such as the maquila-manufacturing in Mexico or the rise of a massive financial sector.

commonly employed inequality indicator, especially in the non-academic discourse (for a comparison of the Palma with a variety of decile and quintile ratios see Palma 2006, 2011).

2. Contextualisation

Concern with the unequal distribution of resources has existed for a long time. However, the concept of what is considered unequal depends on the beholder and her (historical) context. For instance, Plato (427–347 BC) thought that any disparity greater than a 5:1 income ratio between the rich and poor Athenians would trigger ethical imbalance and lead to social ills; George Orwell claimed that an income ratio between the top and bottom of the population's pay above 10:1 was the threshold; and the US financier, John Pierpont Morgan, held that anything beyond a ratio of 20:1 would be inappropriate (Groom, 2014). Nonetheless, in 2013, the actual pay difference between the top Financial Times Stock Exchange companies' (the "FTSE 100") chief executives' and their employees' average pay lay at around 120 times,⁴ while that of some multinational companies' chief executive officers (CEOs) reached 427 or 653 times their employees' median pay. How do we accept these levels? Has our preference for or acceptance of inequality changed that dramatically, or are we just ill-informed about the actual levels of inequality and the harm it is inflicting on our societies? While the first option could be partially responsible for the current circumstances, it is likely that the latter plays a large role as well, i.e. we are not aware of the full extent of today's inequality.

This is at least partially because the conventional inequality measure of the Gini coefficient, an index number between 0 and 1, is difficult to interpret for the lay person without training in comparative inequality analysis. The enigma of an inequality level of 42 lowers the sense of urgency to act.

⁴ Up from 47 times in 1998, though down from a peak of 151 times in 2007.

The mainstream economic literature has long viewed inequality as an indispensable, not necessarily negative, side-effect to the development process (Justino et al., 2003), since it was assumed to recede automatically once the country's income level increased. This conviction largely built on the observation of an inverted-U shape trajectory of inequality dynamics along a country's development path that Simon Kuznets made in the 1950s when analyzing tax return data from the U.S. for the early decades of the 20th Century. Many researchers have challenged this dominant belief, suggesting instead that there is a negative correlation between growth and inequality (Bourguignon, 2004; ECLAC, 2010; Palma, 2011; Piketty, 2014).

There is mounting evidence of the harmful impact that high inequality has on a number of aspects of societal well-being. These include the adverse effects it can have on economic development through its association with social and economic exclusion and instability (Galbraith, 2012), imperilling poverty reduction social cohesion and governability by leading to increases in crime, mass migration, and other forms of social and political conflict and instability. The latter then creates insecurity and distrust amongst economic agents, which could jeopardise economic growth and social development (Justino et al., 2003). Moreover, the complex relationship between inequality and growth (and societal development) acts through a great number of diverse channels, making the disentangling of causal factors extremely difficult. Besides the ethical arguments against inequality, these instrumental reasons speak of an economic case for improving the distribution.⁵

The recent availability of new and better datasets on inequality has sparked renewed interest in the topic (the interest in Piketty's 2014 book *Capital in the Twenty-First Century* illustrates this beautifully). The datasets show that inequality is much greater than it was

⁵ Political philosophy departs from the concept of duty of justice, according to which the responsibilities individuals feel towards their important relations for sharing the same country, ethnic group, etc. impose rules of conduct and responsibility; the cosmopolitan view stresses that we must feel responsible for everybody just for being humans.

assumed to be and, contrary to the belief that inequality was a problem affecting only poor, developing countries, increasingly it is the rich countries that suffer its consequences as well.

This renewed interest in distributional dynamics returns to the forefront two related controversies about, on the one hand, the efficient level of inequality and adequate policies for potential redistribution, and on the other hand, on a more basic level, the actual trends and directions of inequality. In other words, and coming back to the above question: how much inequality do we have, how much do we want to have, and how do we get from the one to the other? In this paper I concentrate chiefly on the first part of this question, revisiting world income distribution and discussing meaningful ways to measure it. However, the debates are intrinsically interlinked, since the way inequality is measured has profound implications for policy-making. Let me illustrate briefly how the measurement question transcends a mere ‘technical’ level.

Palma (2006; 2011) argued on the basis of empirical data that inequality levels are defined by changes in the share of total income held by the extremes of the distribution (i.e. top and bottom) only, whereas the middle-income groups of a population are rather homogenous across countries. Accordingly, Palma (2006; 2011) identifies three population ‘groups’ that are subject to distinct dynamics depending on their income level: (i) the bottom 40 per cent, which in most countries live in continuous poverty in spite of (massive) poverty alleviation programmes, unable to increase their share of the total income; (ii) the middle 50 per cent that defend theirs without being able to increase it significantly though, despite efforts to increase their (human) capital via education, etc.; and (iii) the top 10 per cent, who are drifting away from the rest in stratospheric heights because their share of total income is large and increasing.

The implications these findings bear for policy responses in high-inequality countries are wide-reaching. Globalisation has been thought to lead to convergence of incomes, which

would make interference with distributional levels unnecessary.⁶ Nevertheless, the empirical data does not confirm this prediction. Middle-income countries highly interconnected in terms of international trade, ‘graduate’ to the high-income club being still highly unequal. In the Latin America of the 1990s, income inequality was growing as per capita income increased, moving, therefore, in the opposite direction predicted by the trade-related factor-price-equalisation theorem (Palma, 2011). Moreover, middle-income countries with comparable gross domestic product (GDP) per capita levels (For example, Tunisia and Ecuador) now display a wide array of distributional outcomes, undermining the conclusion of a predictable inverted-U shaped pattern à la Kuznets.

Is it then ‘too much’ if, looking at world inequality, the income of the poorest 5 per cent of United States (US) Americans situates them at the 68th percentile of the world income distribution? Where should we draw the line if, while India's richest ventile⁷ does not even reach the poorest US percentiles, Brazilian rich are as rich as those in the US, and their poor are as poor as those in India or China? This overlap of global, between- and within-country inequalities, illustrated by Milanovic (2011), starkly contrasts with the picture of income convergence, showing instead that inequality in the world is alive and well.⁸ If the consequent inequality is accepted as a threat, laissez-faire is out of the question, now that evidence of inequality increasing, if left to itself, is ubiquitous (Piketty, 2014).

Alternatively, proponents of the human capital accumulation theory have long argued for increases in education as a weapon against inequality. However, differences and improvements in educational outcomes are unable to explain differences in inequality levels.

⁶ The Stolper-Samuelson theorem postulates that a rapid increase in international economic integration should have a positive effect on both inequality within and between countries.

⁷ Dividing the total population into 20 equally sized groups, the richest ventile comprises the five per cent top income earners of a country.

⁸ As opposed to Kuznets-like expectations of receding inequality levels above certain income thresholds, it is the rich countries that are currently experiencing a backlash in their inequality levels, as if emulating the notoriously unequal middle-income countries of Latin America. See also Lakner and Milanovic (2013) for an estimate of the global Gini.

Palma (2006) hence challenges these mainstream policy prescriptions as ill-suited to tackle the inequality challenge:

the top income decile is made up of individuals with relatively advanced levels of education, while those in the bottom four deciles have relatively low levels of formal education [...]. So why do these two relatively homogeneously educated groups have such great distributional diversity? In turn, if significant educational diversity is found among the population in deciles 5 to 9—e.g., in terms of the share of the population with secondary and (especially) tertiary education—why does one find such extraordinary similarity in the shares of national income of this educationally highly heterogeneous group? (Palma, 2006, p.11)

Hence, the discussion about adequate means to deal with inequalities is reopened. However, as the errancy of the human capital theory exemplifies, in order to design an ‘adequate’ policy response, it is necessary to know where exactly inequalities lie along the income distribution. But also, what would constitute a threshold to take action in the first place. Those means applied without paying attention to where inequalities are perpetuated along the income distribution increase their chances of failing in the attempt to improve it, for addressing the ‘wrong’ income groups. The evidence about the high end of the world income distribution presented here is to support the search for an adequate measure that might indicate a threshold of inequality signalling a ‘too much’, or provide a benchmark against which to measure ‘success’ in decreasing inequality levels.

A look inside the distributional dynamics also cements the need for a complementary measure to the Gini, which is unable to detect such distribution-internal dynamics as the ‘stable middle’ and which lacks sensitivity in the extremes. As an example, Figure 1 shows the development of inequality in Mexico, a notoriously unequal country, over the past 60 years according to the Gini coefficient, compared to how the Palma interprets it. The figure shows that where the Gini might have us thinking that nothing actually happens, large distributional changes can occur within the distribution, that are picked up upon by the Palma. The capacity of the Palma to better account for extreme inequalities makes it especially pertinent for high-inequality countries (such as Mexico). Its current proliferation as

an alternative to commonly used indicators is therefore to be welcomed (for a discussion about its advantages over the Gini, see Cobham and Sumner, 2013a).⁹

In sum, far from being an objective ‘technicality’, the general usage of the Gini indicator as an inequality measure obscures developments in the high and low ends of the income distributions, where most of these changes occur. Moreover, this focus exacerbates the image of inequality as a static, unchangeable fact (which it is not). Drawing attention to the location of these changes within the income distribution is important for successful redistribution, as well as to increase people’s sense of urgency to act upon inequality. If policies address those parts of the population where changes are least likely to occur, as happens with the human capital theory addressing the middle groups, the effect on inequality is likely to be minor. Hence, the choice of inequality indicator can be decisive for studying the outcome of how high inequality is. Each option entails a value statement based on the underlying assumptions the measure is constructed upon: do we care mostly about changes in the middle of the distribution, or in the extremes? Transfers between which groups of the population? As indicated at the outset, most people worry about (and underestimate¹⁰) concentration of incomes at the top, advocating for further scrutiny of that part of the population, and an indicator able to display this. The introduction of the 10/40 ratio is a decisive change in this regard. I will show in continuation that it could be complemented by another ratio, even more sensitive to what happens at the very top. Following Palma’s lead, it shall be denominated here as the ‘5/40 ratio’.

⁹ Some international institutions, such as the UN economic commissions in Latin America and Asia, and also Oxfam, have adopted the Palma into their portfolio already.

¹⁰ Interestingly, in a 2011 poll about inequality perceptions in the US, where people were asked to place themselves in the bottom 99 per cent or top 1 per cent of the income distribution, 68 per cent of respondents said they were part of the 99 per cent, while 13 per cent thought of themselves as part of the top 1 per cent of money makers (19 per cent had no opinion on the question). This strikingly indicates common misperceptions, both about the actual level of inequality, and the relative position of the different income groups (Jordan 2011).

3. Methodological Concerns

The world's relative distribution of disposable income¹¹ can be scrutinised by looking at the shares of total national income within any country that a subgroup disposes of. For the purpose at hand, these data had to be disaggregated for population ventiles (20 equal-size population groups) or higher. Two main datasets, each assembled from national household survey data, fulfill this condition.

3.1. The Datasets

The first is the World Income Distribution Database (WYD) compiled by Milanovic (2014). It combines data for developing countries from the World Bank's PovcalNet database with LIS (formerly Luxembourg Income Study) data for the rich countries.¹² This dataset covers about 90 per cent of world population, and 93 per cent of its GDP in 2008 (Lakner and Milanovic, 2013), justifying generalisations about the world income distribution. The WYD-2008 dataset comprises data on household per capita income (or expenditure) for individuals, where all household members are assigned the same (average) income/expenditure, i.e. any economies of scale in household consumption and within-household inequality are ignored (Milanovic, 2009).¹³ While some earlier versions of this dataset disaggregate the population by ventiles, only the benchmark year of 2008, the latest year available, publishes all 116¹⁴ countries' income according to ventiles, thus allowing for the calculation of the top 5 per cent's income share. Consequently, the findings from this dataset inform the issue of

11 Relative income distribution as a proxy for inequality has the benefits over most other measures in that it is relatively easily comparable across programmes, policies, time, and geographic space, as well as reasonably well-documented.

12 See also Milanovic (1999) for details on the sources for the surveys and (Lakner and Milanovic, 2013) for further details on the construction of the dataset.

13 This is because per capita incomes have the advantage that they are simple to compute and have natural counterparts in the national accounts (which do not compute equivalent incomes) (Lakner and Milanovic, 2013).

14 In 2008, Switzerland's bottom decile had an average income of zero. The WYD reports this as missing. It is therefore not included in this sample, since the bottom 40 per cent share cannot be calculated without the first decile.

distributional dynamics when the income share of the richest 5 per cent income earners is considered, and how this information could help measure inequality.

The second dataset is the LIS (2014)¹⁵ household survey microdata for a subset of 41 upper- and upper-middle-income countries reporting shares of household disposable income for the top percentile. Here, this source is utilised for the investigation of the richest 1 per cent, as well as the comparison over time. No top-coding is applied, but household members' incomes are weighted according to the equivalent adult scheme. Since LIS data is harmonised to allow for time-series comparisons,¹⁶ although much smaller than the WYD-2008, this sample hosts the largest compilation of income distribution data for developments over time for the whole population, yet disaggregated up to the top 5 or 1 per cent. Since surveys do not exist for all years and countries, and to maximise the sample of countries for the subsequent comparison over time, two year-points are chosen, the first point ('1990') includes all countries that present a survey in the early 1990s, while the later point ('2010') essentially includes all surveys from the late 2000s and 2011.¹⁷

3.2. Data limitations

Consciousness about the limitations of available data is opportune. While household survey data is no perfect source for comparative studies over time and/or space (diversity in methods applied for data collection; variation in choice of indicators and measures used, e.g. between income and expenditure; changing periodicity of surveys), it remains a good starting point for a thorough investigation of world inequality patterns. However, while the quality of published datasets has improved, availability is still limited and comprehensiveness increased

¹⁵Considering that approaches to harmonisation differ between the two datasets, the resultant estimates for countries may differ accordingly. These discrepancies are small though and do not affect the results of the investigation.

¹⁶ For details on the harmonisation of the microdata see LIS' webpage at <http://www.lisdatacenter.org/our-data/lis-database/documentation/> (last accessed at 20.03.2015).

¹⁷ A minimum of 15 years is guarded between the two data points, except for the case of Australia, Austria, and Slovenia, where the difference is 8, 11, and 14 years, respectively.

only very recently, when attempts have been made to improve comparability between different countries and over time. Nonetheless, due to the lack of appropriate historical data, time-series testing is difficult for larger samples (Palma, 2011). Besides, there are intrinsic problems with household survey data when it comes to reporting top incomes. The difficulty of reaching these households (considering that they are rare and often difficult to access) creates sampling issues leading to the understatement of incomes at the top (Lakner and Milanovic, 2013).

However, at least three reasons speak for using the available data anyway. Firstly, the problem of under-reporting is intrinsic to all self-reported information, and likely to affect (to varying degrees) any kind of data we are able to collect on the income distribution, and at all levels of income. Alternative empirical approaches try to escape these shortcomings through usage of industrial data (Galbraith, 2012) but suffer from additional severe limitations, including the difficulty to account for informal labour, rendering them less suitable for the purpose at hand.¹⁸ Other types of research, such as Piketty's and Kuznets' that rely on tax accounts, might appear to be more accurate at measuring top incomes. While arguably more reliable as to their objectivity (giving less leeway to individuals' disclosing income at their own discretion), they are constrained by the same informal sector problem, as well as by massive under-reporting and tax evasion, especially in developing countries.¹⁹ Besides, it is quite likely that especially the highest income groups will have the knowledge and capacity to deduct large sums from their incomes ('income minimisation', a problem likely to be particularly acute in developing countries too; see Lakner and Milanovic (2013)) exposing

¹⁸ Reliance on manufacturing data proves highly problematic in a developing country context where often over half of the working population is employed in the informal sector of which few accounts exist, and large shares of the labour force remain in agricultural (or services) sectors with entirely different income structures. Furthermore, it can be argued that manufacturing is not representative for the whole breadth of incomes, again making it difficult to extrapolate to the rest of the population.

¹⁹ In Mexico, for instance, only about 30 per cent of workers pay taxes, restricting the representativeness of the sample and the predictive capacity of such numbers.

this data to a bias as well. Moreover, few countries make these numbers available. For the purpose at hand, even if they were available, they would not be suitable, since they do not report incomes for the lower-income groups.

Secondly, this information is what we have. Pragmatically, just because we do not have better data so far does not mean we should completely dismiss those we have. While critical awareness of their shortcomings is necessary, they can still tell us important stories that might further our knowledge about the state of the world we live in.²⁰ Thus, until, for instance, efforts were to be made to combine tax report and household survey data to represent both the lower- and upper-income groups more adequately, household surveys seem to be the most reliable source we have for the purpose at hand. In this sense, the updated WYD-2008 dataset by Milanovic used in this study, which includes data by ventiles for a large part of the world, is an important improvement.

Thirdly, considering that even the limited data we have, with all its understating, shows such stark income polarisation, it is likely that the ‘real’ distribution is even worse, only accentuating the argument. Nonetheless, it is bearing all of the above shortcomings in mind that the conclusions proposed here should be approached.

4. **Analysis**

Not only is there a high concentration of income in the shares of the very top percentiles of the distribution, but concern over inequality levels would presumably be higher if its actual dimension were more readily accessible. To create such awareness, the measuring indicator must be able to make high levels of concentration explicit. Palma (2006) suggested the usage of the ratio now known as ‘the Palma’, whose key feature is its focus on the tails of

²⁰ Atkinson (1983, p.206) quotes Kuznets: ‘the approach adopted here “is needed not only to warn against uncritical reliance on the estimates as they are usually shown but, perhaps even more, to counter their complete dismissal because of the serious errors and biases in them”.’

the distribution where change is happening, particularly at the top decile.²¹ This indicator thus comes much closer to describing actual income dynamics (compared, for instance, to the Gini). However, it is very likely that the distribution within this decile is highly unequal as well. So what happens if we disaggregate the top decile and look at the top 5 per cent instead? How does its share compare between countries and over time?

4.1. Looking Inside the Top Decile

Using the WYD-2008 dataset, let us look at the comparison between countries first. Panel A in Figure 2 shows the cross-sectional comparison of the 116 countries in the sample, ranked according to the share of total income the lowest income groups of their populations' hold.

Two features are pervasive: firstly, Palma's stable middle is easy to appreciate; across the sample, the middle 50 per cent of the population hold on average 52 per cent of total income, while differences in the shares of bottom and top deciles are large. This contrast between the homogenous middle and the heterogeneous extremes of the income distribution is reproduced in Table 1.

Before turning to that, though, Figure 2 has a second salient feature: that the second highest ventile is extremely stable across the countries in the sample, and that differences in the top decile are defined almost entirely by the size of the highest ventile. Just to visualise this homogeneity in the 19th ventile, Panel B inverts the image, aligned now according to the size of the richest 5 per cent.

This reversed view emphasises how in those countries where the last ventile is extraordinarily large, the 'extra' part cannot be taken from the low-income groups anymore, and hence the share of the middle in these countries is somewhat smaller than that of the rest.

²¹ Since other advantages of the Palma over the Gini and other traditional measures of inequality have been analysed in Cobham and Sumner (2013a; 2013b) and Doyle and Stiglitz (2014), here the focus shall lie on the complementary features of related ratios to the Palma rather than comparing them to the Gini again.

But this extra part is not distributed equally within the top decile—the benefits accrue almost exclusively to the top 5 per cent, while the second-to-last ventile increases only marginally. For most of the other countries there is a direct relationship between the bottom 40 per cent’s share increasing where the top 5 per cent’s decreases, while the middle 50 per cent and the second-to-highest ventile remain overall fairly stable. Differences in the level secured by the top ventile in the respective countries are vast: shares here range from 11.5 per cent (in Slovenia) to almost 39 per cent of total income (in South Africa)—more than three times the size. Analogously, the bottom 40 per cent of the population hold on to a mere 5.5 per cent at the low end (South Africa) compared to a maximum of almost 5 times this share (25.5 per cent in Slovenia). These large differences contrast sharply with the relative homogeneity in the upper-middle-income groups, where the augmented upper-middle 55 per cent differs less than 20 per cent between the highest (China) and lowest (Central African Republic) share.

Table 1 shows that the coefficient of variation for such an augmented middle (including arguably very upper-middle 19th ventile) is very small; such a population share would be even more homogenous than the middle 50 per cent. Variations of the bottom and top groups, on the other hand, are a lot larger—especially that of the highest ventile. These observations are in line with Palma’s proposition according to which differences in inequality levels in the world are explained overwhelmingly by the differences in the shares of the extremes, while the middle-income groups own a similar share of total income all over the world. Palma found this relationship to hold when using the top decile as the delimitation of the elite income group. The new data analysed here shows that this relationship holds more, i.e. the middle becomes more homogeneous, when a smaller fraction at the top is considered (top 5 per cent), and analogously the middle is ‘augmented’ to encompass 55 per cent of the population.

If it is the case that the most stable sub-part of the population is the augmented middle, while the most diverse part is the top percentiles, such findings would warrant the usage of a sort of ‘higher order’ Palma ratio—a ‘5/40 ratio’. For such an indicator, smaller shares of the elite’s income, For example, the highest 5 (or 1) per cent, are divided over that of the poorest 40 per cent of the population.

What speaks for the usage of a 5/40 ratio, as a complementary measure to the Palma, is the extraordinary concentration of income within the top decile, which in itself is much more unequal than any other decile. Whereas the 90th to 95th percentiles receive only somewhat more than the 18th ventile, which lies just below it (20 per cent more on average), for all 116 countries, the last ventile secures on average almost double compared to the one preceding it (10.3 per cent of total income for the 19th ventile and 19.6 per cent for the last one). For 34 countries, its share is more than double, and three countries—Guatemala, Central African Republic, and Honduras—more than triple that of the 5 per cent just below. No other ventile is marked by such drastic difference between the ones just next to it (whether that is above or below). Figure 3 below illustrates the relationship between the 20th ventile and the tenth decile, and the 19th ventile, respectively. As can be observed, the trajectory of the top 10 per cent and the top 5 per cent is not only in close alignment, but also very diverse across countries. Meanwhile, the 19th ventile remains remarkably constant across countries, and does not display the same shape of large variance between countries.

On the basis of these findings, we can construct a 5/40 ratio, which, on a world average, takes the value of 1. This is a very interesting feature, since it shows the unity of the top 5 per cent (keep in mind that this is not the global 5/40 ratio though, but a country average) defending the same share of total income as the 40 per cent with the lowest incomes. This translates into a person in the top ventile owning eight times the share of a person in the bottom 40 per cent of the population. Moving away from the average, at the low end this

difference shrinks to a 5/40 ratio of 0.45 times in Slovenia, whereas at the other extreme, in South Africa, this indicator reads just over 7, meaning that a person in the top 5 per cent owns on average 56 times the share of total income a person in the bottom 40 per cent can hope to have! Such absurd inequality seems to be a unique trademark of some Southern African and the Latin American countries. The latter region, for instance, features a 5/40 ratio of 2.7 when isolated from the rest. Of the top 20 most unequal countries within the sample, 15 are Latin American, and only Uruguay, Argentina, and Venezuela are located further down (rank 26, 27, and 34 respectively²²).

The 5/40 ratio shares characteristics with the Palma: The most stylized features of the Palma, which cannot be observed by the Gini-inequality ranking, are, in Palma's (2014, p.4) words, that inequality across the world, as measured by this ratio, "increases relatively slowly at first, and almost linearly, only to switch gear when Latin American countries enter *en masse* (around ranking 100) – and then increases rapidly and geometrically." Palma (2014) continues to assert that, had the 'steady pace' of the first four-fifths continued, the maximum Palma ratio should be around 3, rather than triple this amount. A similar pattern can be observed with the 5/40 ratio, which is slightly more linear in its lower rankings, and increases even more steeply once 'the usual suspects' enter the ranking (shortly after rank 100) (see Figure 4). Data points for South Africa are off the chart, literally (with a Palma of 9.9 and a 5/40 ratio of just over 7) – is this 'too much'?

Just like the Palma, a 5/40 ratio would lie fairly close to the Gini at levels of lower inequality. However, in countries with higher inequality, the difference between the two (three) indices increases dramatically, testifying to the higher sensitivity of the Palma and 5/40 ratio to highly polarised distributions. Figure 4 also shows that the rankings ordered by the Palma do not in all cases coincide with those of the 5/40 ratio: Guatemala, Uruguay, or

²² Considering high inflation in Argentina and Venezuela, there might exist measurement challenges in these two countries though (Palma, 2014).

Uganda for instance rank higher in the 5/40 ratio than their corresponding position along the Palma scale would have us assume. Israel or Thailand, on the other hand, display a lower 5/40 ratio ranking than their Palma implies. This shows that the concentrations at the very top are not necessarily predictable by the look of their top decile, increasing thus the two measures' complementarity.

Summing up, the 5/40 ratio could be a useful supplement to the Palma, offering some unique features beyond the attributes of the Palma (of which many are shared by the 5/40 ratio, such as its intuitiveness and simplicity, making it easy-to-interpret, and suited for policy-making). It is even more sensitive to changes in the upmost ranks of the income distribution, while accounting for the fact that the top decile in itself is highly unequally distributed. The elegance of the 50/50 rule makes the Palma uniquely suited for policy purposes; however, the 5/40 ratio's unity-relation make the latter a good ally for measuring what most people care about when inequality is the concern: the concentration of income at the very top, at the expense of the low-income groups. But how do these traits feature when the top is disaggregated even further?

4.2. Disaggregating the Top Ventile

Unfortunately, comparable data for the top 1 per cent of the income distribution is even sparser. The LIS database hosts 41 countries for which such numbers are available for at least one year. Figure 5 displays their richest populations' share of income and how this ranks them relative to the remainder of countries. It thereby illustrates the importance of dissecting the top decile. While the top rankings are no surprise to a scholar of inequality, some of the middle positions run counter to expectations based on Gini rankings (or even Palma).

Considering the data limitations introduced above, results are indicative rather than exact. Notwithstanding, the figure shows that scoring relatively well in terms of inequality according to measures such as the Gini does not necessarily coincide with less concentration

of incomes at the very top end of the distribution, as the cases of Belgium or Denmark (Gini of 26.2 and 25.2, respectively, in 2010) illustrate.

While five countries (Colombia, Guatemala, Peru, South Africa, and India) compete for the top positions amongst themselves and lead the table in all three disaggregated groups, beyond these top five, the lines become more blurred. The strong regional dynamics discernible when analysing the top 10 per cent of the income distribution are less marked when looking at the top 5 or even 1 per cent. This is especially so when it comes to those countries perceived as more equal according to other measures of inequality – namely, the Northern European ones, and especially the Nordic countries, where Denmark, Finland, and Norway lie exactly in the middle of the 41 countries presented when measured by their top 1 per cent income holders. Nonetheless, it is also clearly noticeable that all the Latin American countries included remain faithful to their reputation and make the top ten.

Interestingly, some of the Asian countries (Japan and Korea) rank significantly more equal according to their highest percentile than when the 20th ventile or top decile is considered, and the same is true—and even more pronounced—for Spain, which has the second smallest share for the top 1 per cent in the sample (though not when the top 2 or 3 per cent is considered), and Estonia and Ireland. This shows a distributional diversity within the elite, which does not necessarily run along the same lines as the rest of the distribution. Such could either be an indicator of more ‘flat’ structures within the elite in these countries, or a sign that the concepts used do not capture the existing concentration (for example, a 0.1 per cent or smaller share would be necessary to see further divergence). An interesting case warranting further investigation for its consistently low inequality is Slovenia, ranking last (i.e. most equal) in terms of its top 1, 2, 3, 5, and 10 per cent’s share, as well as Palma and Gini (incidentally, the country also features the second highest minimum wage in relative terms in the world).

Combining the information the Palma-family can provide, Figure 6 shows two interesting dynamics comparing the three ‘Palmae’: on the one hand, the ‘original’ Palma seems to capture more movement for the rich, middle-inequality countries compared to its derivatives. On the other hand, especially at higher inequality levels (South Africa is off the chart again), the three indicators lose the fairly strong correlation they display for countries with low inequality.

Clearly, the set of countries for which data on the top percentiles is available is too small and regionally biased (it cannot account for most of Africa and large parts of Asia) to make definite statements about this issue on the world scale. However, considering that it comprises a number of the notoriously most unequal countries, as well as the most equal ones in terms of most indicators, it allows for the formulation of some tentative conclusions. In this sense, the same phenomenon observed above, describing the difference between the 19th and 20th ventile, can also be identified when comparing the 100th and the 99th percentile of income: in many cases the former is more than double the size of its closest neighbour, this distance being far larger than that to the next-lower, 98th, percentile.

While it seems difficult to change the fact of high concentration at the very top, apparently some countries are nonetheless able to significantly enlarge their bottom part’s share. For instance, even though the Scandinavian countries still maintain a large top 1 per cent share, their bottom deciles appropriate approximately 23–25 per cent of total income—compared to the lowest four deciles of the Latin American countries getting only about half of that. That this is not an intrinsic feature of their income distribution but the result of conscious policy-making becomes clear when comparing the pre-tax and transfer distribution of the Nordic with that of the Latin American countries.

Unfortunately, the information about pre-tax and transfer income distributions is even more limited than that of disposable income, and difficult to compare. Hence, a direct

comparison between two countries (based on simulations), rather than a larger sample, must suffice to point to the policy importance in this instance. In the example presented in Table 2, the level of inequality according to market income (as measured by the Gini) in Mexico and Denmark is not that different, but the latter manages to reduce its disposable income Gini to almost half its size, whereas the Mexican Gini remains almost intact. While we cannot compare the market income of the top percentiles to their disposable income due to the lack of available data, it seems that, unsurprisingly, redistributive policies in Denmark have a larger equalising impact than those in Mexico, even if this leaves the hegemony of the top 1 per cent untouched.

Such comparison shows that the high inequality levels observed in certain countries are not unchangeable in principle. With the right policy application, they could be diminished significantly—at least insofar as the share of the poor groups could be augmented. However, this table scarcely discloses any information about the concentration at the top. As Figures 5 and 6 showed, depending on the choice of indicator, inequality rankings will be altered. This alludes to the importance of choosing the ‘right’ indicator for the purpose at hand.

However, as we just saw when disaggregating the very top, the highest percentile in Denmark is also significantly higher than we would expect from the size of its top decile, or ventile even (and certainly compared to its Gini). Nonetheless, with 23.4 per cent of total income, its bottom 40 per cent of income earners secure one of the largest shares for their group in the sample. At the same time, its middle groups (deciles 5 to 9) own more than 55 per cent of total income, placing it above the sample’s average. This must mean that the country manages to redistribute from its top decile towards the bottom, yet leaving its highest percentile as well as middle groups relatively intact. Accordingly, a 1/40 ratio would place Denmark exactly at the middle of our sample, while the Gini, Palma and even 5/40 would place it in the top ten of most equal countries. Were there data available for smaller fractions

of the elite (0.1, 0.01 per cent etc.), it is conceivable that rankings would look yet different again.

It becomes clear from the empirical data presented that the different inequality measures rank countries differently. While in itself this is not a revolutionary finding, it does mean that the undifferentiated usage of the Gini coefficient to describe and compare inequality levels among countries has serious shortcomings as to its accuracy when the distributional extremes are scrutinised. However, indeed any one (summary) measure will potentially face problems when attempting to describe every change in a distribution where different forces are at work at the same time and prioritise diverse aspects of inequality. Depending on the stated objective of the investigation it then becomes paramount to choose the adequate corresponding indicator(s). The impact of this choice is accentuated in high inequality countries, where a similar Gini coefficient compared to another country can mask large differences that are captured by a Palma ratio or its derivatives instead: for instance, comparing Portugal and Sierra Leone, both feature a Gini of 34.4 in 2010—however, their Palmas of 1.38 and 1.73, respectively, position them more than 20 ranks apart from each other on a global scale.

4.3. 5/40 or 1/40 ratio?

Taking into account the vast inequality within the top decile of the distribution, it makes sense to include a 5/40 or 1/40 ratio measure into the portfolio of available indicators. This is particularly relevant in high-inequality countries, where the lower-income groups of the population feature a particularly small share of total income combined with the high concentration at the top end. Just like the Palma, the derivatives are more intuitive than the Gini in its interpretation ('the top 1 per cent earns x times what the bottom 40 per cent receive'), and simpler in its calculation. The 1/40 ratio's disadvantage compared to the other two is that its values mostly lie below 1, making interpretation slightly more cumbersome and

hence less practical to manage ('in Brazil, the top 1 per cent's share of total income is 0.635 times that of the bottom 40 per cent'); it lacks the elegant feature of the Palma of offering a simple-to-handle number as a result. Notwithstanding, it could be a useful supplement for the more traditional indicators and the Palma. However, between the two concentration measures presented here, the 5/40 ratio seems more suitable than the third order Palma, not only because the top 1 per cent of income earners are hard to come by with household surveys (and most attempts to measure them anyway) and only few countries make related data available, but also because only three countries in the sample feature a 1/40 ratio above unity,²³ making interpretation of the resulting fractions less elegant.

Therefore, the unique feature making the 5/40 ratio interesting in this context is its potential to function as a threshold of sorts. As long as its value lies under 1, numerical differences are slight. However, as soon as the threshold of 1 is surpassed, its values increase very quickly. As stated earlier, world average for the 5/40 ratio is 1. Of the 116 countries in the sample, half have a 5/40 ratio above 1 (57 countries), some significantly so (the Latin American and some Southern African countries, up to 7 in South Africa). In this sense, a 5/40 ratio of 1 could be a signalling limit for excessive inequality: when the top 5 per cent of income earners keep a larger share of total income than the bottom 40 per cent, it is 'too much' and alarming inequality and concentration levels have been reached. Subsequent thresholds might indicate the urgency with which this excessive inequality must be addressed (For example, 5/40 ratio above 2: prioritise; above 3: top priority.) While countries might have their idiosyncratic tolerance for inequality, the 5/40 ratio=1 lends itself as an upper bound to be aspired to -- for example, as an objective for the forthcoming Sustainable Development Goal indicators -- besides being an intuitive and 'catchy' relationship to

²³ Information about this fraction of the income distribution is only available for a small subset of countries, so no generalisations about the world distribution can be made on the basis of the available evidence, especially considering that this subset is comprised mainly by the rich OECD countries, of which at least the European subset (especially the Northern European ones) are less unequal than the world average.

remember also for lay persons. Considering the concentration of income at the top of even the highest decile, it seems useful to combine such a measure with the Palma ratio.²⁴

It should be clear by now that inequality differences are largely decided in the extremes of the distribution, highly concentrated at the smallest fractions of the top. Is this a pattern reinforced over time, i.e. have the tails diverged over the last decades while the middle's share has simultaneously become consolidated?

4.4. **Developments within the Top Income Shares Over Time**

That Lakner and Milanovic (2013) find 44 per cent of the increase of global income between 1988 and 2008 to go to the top 5 per cent of world population points towards a divergence of the tails. The four developing countries that are represented in the WYD dataset by ventiles from 1988–2008 (Colombia, Pakistan, Philippines, and Thailand) testify to these trends, as Figure 7 illustrates. Although this sample of four countries is obviously too small to make any generalisations, its income distribution develops exactly as we would expect from our hypotheses: the middle groups remain relatively stable (their share drops slightly from an average of 50.6 per cent in 1988 to 50 per cent in 2008), while the top increases its share mostly at the expense of the bottom (poorest four deciles see their share shrink from 16 per cent in 1988 to 14.7 per cent in 2008), shifting the gain mostly onto the highest ventile (from 22.3 per cent in 1988 to 23.8 per cent in 2008). The subsample's 5/40 ratio increases from 1.5 to 1.93 over the period considered. Their 19th ventile, on the other hand, remains extremely stable, confirming that the increase at the top comes at the expense of the bottom.

Do such dynamics hold for the top 1 per cent as well? To test a distributional trend over the last two decades, a subset of 25 countries for which disaggregated data exists for the

²⁴ From the WYD dataset, only 14 countries have a Palma below 1, making it an ambitious threshold to aspire to at the outset. However, considering the different concepts the two indicators measure, a combination of both Palma and 5/40 ratio below 1 as threshold for excessive inequality could be envisioned.

top 5 and top 1 per cent for the early 1990s and the late 2000s is selected from the LIS dataset.²⁵ Although there are important exceptions, in most countries the income share of the rich increased for all three concepts of the elite (the top 10, 5, and 1 per cent of income earners).

Table 3 shows that, on average, the bottom deciles lost a significant percentage of their share over the period under consideration, while the largest corresponding increase accrued to the top decile. Within this highest income elite, the benefits clearly fell to the 20th ventile (0.7 per cent increase in their share of total income), as compared to the 19th (0.2 per cent increase). Meanwhile, the middle and upper-middle groups remained virtually unchanged (albeit a slight increase can be detected in the augmented middle 55 per cent). Besides, two interesting tendencies can be observed: firstly, the coefficients of variation are much smaller for the middle-income groups' shares than for the extremes (top and bottom), both in 1990 and in 2010, especially if the top percentiles are taken separately (i.e. top 5 and 1 per cent). Secondly, however, all of these coefficients decrease over the period considered (all ranges become smaller), where particularly large decreases can be observed for the deciles 5 to 9 and top 5 per cent groups.²⁶

Moving away from the averages into the individual countries' dynamics, the vast majority of countries saw their elites' shares increase; the same seven countries that experienced a decrease in their top 5 per cent also saw their top 1 per cent shrinking (with the largest decreases observable in Mexico and Ireland). Only in three of the seven countries displaying a decrease in the top 1 per cent does this exceed 1 per cent though (Mexico with

25 This small dataset unfortunately does not include any developing countries, making it less representative on a world level. However, it is the only comparable dataset including information about the respective countries' entire income distribution, and separately listing the top 1 and top 5 per cent.

26 However, the coefficients of variation are rather sensitive to changes in the sample of the countries (For example, if Mexico is taken out, the middle 50 per cent's coefficient decreases to 0.017 for both years, whereas that for T1 decreases from 0.227 in 1990 to 0.185 in 2010). A different subset might yield different results. Taking a conservative approach, the observed trends might thus be most useful as a starting point for further research, rather than statistically robust results.

3.5 per cent, Ireland with 2.5 per cent, and Switzerland with 2.4 per cent, see Figure 8). As to changes in the middle groups, half of the countries in the sample experienced slight increase, the other half slight decrease; in two-thirds of the countries these changes were less than 1 percentage point of total income, and except for changes in Mexico and Spain (5.3 and 2.5 per cent increase for the middle deciles, respectively), the remaining six countries saw changes of less than 2 per cent. Developments in the low end of the distribution are less encouraging, although moving in the direction predicted by Palma's hypothesis: in two-thirds of the countries the share held by the poorest four deciles shrank, and of those countries where it increased, only in three countries did this positive change surpass the 1 per cent mark (Greece and Mexico, with 1.1 per cent increase each, and Ireland with 2 per cent). This general picture shows that in the majority of countries the bottom groups lost income share while the top groups gained ground, whereas in the middle part no clear tendency can be discerned for this subset of countries.

It seems, thus, that indeed the middle and upper-middle groups are more homogeneous now than two decades ago. However, even the top and bottom parts in our sample are 'converging' according to their diminished spread (if less so). At the same time, the tails seem to be drifting apart, albeit consolidating their shares. Further research would have to investigate the decisive factors behind these patterns, but it is likely that declining top tax rates and capture of the institutions via 'rigged political rules' (Oxfam, 2014) by the elites, as well as stagnating real wages in many of the countries considered, play a role.

5. Conclusion

The findings presented here confirm Palma's predictions of a homogenising middle and diverging tails within the income distribution. But they also show that it is necessary to look into the top decile to see the 'real' income distribution. While the share of the bottom 40 per

cent of income earners decreased over the time period considered, that of the top shares, whether one looks at the top 1, 5, or top 10 per cent, increased—i.e. the tails drifted apart. Meanwhile, the middle- and upper-middle-income groups (deciles 5 to 9, augmented middle 55 per cent, upper-middle 35 per cent, even percentiles 90 to 95 to some degree) not only remained stable but homogenised insofar as their spread decreased (with deciles 5 to 9 and the augmented middle 55 per cent showing the lowest internal divergence, and top 1 per cent the highest, followed by the 20th ventile, for both years—just as expected).

This concentration at the very top speaks for the usage of an indicator that can account for this pattern. The empirical evidence points to a lack of fit into either the Gini-world of a broad variety of incomes experiencing changes around the middle of the distribution, or the simple concentration of incomes at the top of the distribution—rather, it seems to indicate two different, co-existing worlds within each distribution: one of decreasing inequality within the middle groups, and one where the poor and rich continue to drift apart (with the top percentiles living in a world of their own—and surely, if we could look into the top 1 per cent, distances would increase even more). Which one do we want to live in? Do we want the kind of inequality we have? If not, how do we change it?

Two essential findings rise above all others: i) the fact that, in every country, the top 5 per cent's income share is so much larger than the next five per cent's (i.e. the distribution in the top decile itself is very unequal), at times exceeding the triple value of the following ventile; and ii) in the struggle over the remaining income, where the top 5 (rather than 10) per cent manages to increase their share, it is usually at the expense of the bottom 40 per cent, rather than the middle 50 per cent, whose share remains fairly stable at just over 50 per cent of total income. Only where the bottom is so squeezed already that they would literally be left with zero income, do gains at the top correspond to losses for the middle groups. Such patterns can hardly be considered in the interest of the majority of people, and resemble a world governed,

in Stiglitz's (2011) words, "of the 1 per cent for the 1 per cent by the 1 per cent." It is likely that most people, if aware of these trends, would oppose them.

The currently widespread ignorance of these actual inequality patterns, where extreme wealth and poverty are both underestimated and the 'stable middle' does not feature at all, obstructs the application of effective remedies to improve the distribution. It also means that the choice of inequality indicator, far from being a dry, technical exercise, is indeed highly (politically and ideologically) charged. Policy prescriptions to resolve the inequality crisis have evolved into the (politically cheap) increase of access to education, in combination with (financially cheap) patchy anti-poverty programmes (such as conditional cash transfers), rather than tackling the issue at its head, the high end of the distribution (via fiscal and labour market policies, such as progressive direct taxes, but also minimum wage policies.) Arguably, this has to do with the fact that the actual dimension of inequality is not stated explicitly enough in the indicators we use to describe it.

With better data becoming available to describe the actual distributional patterns, this investigation shows that the reliance on one inequality indicator, be it the Gini or another (however sophisticated) to account for the whole income distribution, is risky considering the opposing tendencies within the distribution. While the Gini, and/or Palma, is not to be substituted, in order to truly display the degree of polarisation in countries' income distribution, they should certainly be complemented by an indicator accounting for the top 5 per cent, such as the '5/40 ratio' proposed here. By combining these indicators, a more realistic picture of the income distribution can be provided.

References

- Atkinson, A.B. (1983). *The Economics of Inequality*. Oxford: Oxford University Press.
- Bourguignon, F. (2004). The Poverty-Growth-Inequality Triangle. *New Delhi Working Paper* No. 125. New Delhi: Indian Council for Research on International Economic Relations.
- Chang, H. (2014). Development Economics Course Lecture Notes, Lent Term. Cambridge: Centre of Development Studies, University of Cambridge.
- Cobham, A. and A. Sumner (2013a). Is It All About the Tails? The Palma Measure of Income Inequality. *Centre for Global Development Working Paper* 343. Washington, DC: Centre for Global Development.
- Cobham, A., and A. Sumner (2013b). Putting the Gini back in the Bottle? The Palma as a Policy-Relevant Measure of Inequality. King's International Development Institute. London: King's College London.
- Doyle, M.W., and J.E. Stiglitz (2014). Eliminating Extreme Inequality: A Sustainable Development Goal, 2015–2030. *Ethics and International Affairs*, 28(1), 5–13.
- ECLAC (2010). *La hora de la igualdad: brechas por cerrar, caminos por abrir*. Santiago de Chile: United Nations.
- Forbes Magazine (2014). The Worlds Billionaires. Available at: <http://www.forbes.com/billionaires/>. Last accessed on 30.10.2014.
- Galbraith, J.K. (2012). *Inequality and Instability*. Oxford: Oxford University Press.
- Goñi, E., J.H. López, and L. Servén (2008). Fiscal Redistribution and Income Inequality in Latin America. *Policy Research Working Paper* No. 4487. Washington, DC: The World Bank.
- Groom, B. (2014). If only Plato were running Barclays. Available at: <http://www.ft.com/cms/s/0/c0ecff02-a84a-11e3-a946-00144feab7de.html#axzz31PUsfSmr>. Last accessed on 04.08.2014.
- Hasenheit, M. (2014). Inequality: Where Palma is better than Gini, Network for Green Economy Indicators. Available at: <http://netgreen-project.eu/blog/2014/06/18/inequality-where-palma-better-gini>. Last accessed on 04.05.2014.
- Justino, P., J. Litchfield, and L. Whitehead (2003). The Impact of Inequality in Latin America. *PRUS Working Paper* No. 21. Sussex: University of Sussex.
- Jordan, E. (2011). Can 13% be part of elite 1%?. Available at: <http://pollposition.com/2011/11/01/can-13-be-part-of-elite-1/> Last accessed on 20.03.2015.

- Krozer, A., and J.C. Moreno-Brid (2014). Inequality in Mexico. *World Economics Association*, 4(5), 4.
- Krueger, A.O. (2003). Address on Globalization at the 7th St. Petersburg International Economic Forum. St. Petersburg. Available at: <https://www.imf.org/external/np/speeches/2003/061803.htm>. Last accessed 20.04.2014
- Lakner, C., and B. Milanovic (2013). Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. *Policy Research Working Paper* No. 6719. Washington, DC: The World Bank.
- Luxembourg Income Study (LIS) Database. <http://www.lisdatacenter.org> (multiple countries; accessed June 2014-March 2015). Luxembourg: LIS.
- Milanovic, B. (2014). World Income Distribution Dataset 2008. Available at: <http://econ.worldbank.org/WBSITE/EXTERNAL/EXTDEC/EXTRESEARCH/0,,contentMDK:22261771~pagePK:64214825~piPK:64214943~theSitePK:469382,00.html>. Last accessed 20.03.2015.
- Milanovic, B. (1999). True World Income Distribution, 1988 and 1993. First Calculations, Based on Household Surveys Alone. *Policy Research Working Paper* No. 2244. Washington, DC: The World Bank.
- Milanovic, B. (2009). Global Inequality and the Global Inequality Extraction Ratio. *Policy Research Working Paper* No. 5044. Washington, DC: The World Bank.
- Milanovic, B. (2011). *The Haves and Have-nots*. New York: Basic Books.
- OECD (2014). Focus on Top Incomes and Taxation in OECD Countries: Was the Crisis a Game Changer?. Available at: <http://www.oecd.org/els/soc/OECD2014-FocusOnTopIncomes.pdf>. Last accessed on 01.08.2014.
- Oxfam (2014). Working for the Few. Political Capture and Economic Inequality. *Oxfam Briefing Paper* No. 178. London: Oxfam.
- Palma, J.G. (2006). Globalizing Inequality: ‘Centrifugal’ and ‘Centripetal’ Forces at Work. *DESA Working Paper* No. 35. New York: United Nations.
- Palma, J.G. (2011). Homogeneous Middles vs. Heterogeneous Tails, and the End of the ‘Inverted-U’: the Share of the Rich is what it’s all about. *Development and Change*, 42(1), 87-153.
- Palma, J.G. (2014). Has the income share of the middle and upper-middle been stable over time, or is its current homogeneity across the world the outcome of a process of convergence? The ‘Palma Ratio’ revisited. *Cambridge Working Papers in Economics (CWPE)* No. 1437. Cambridge: University of Cambridge.
- Piketty, T. (2014). *Capital in the Twenty-First Century*. Harvard: Harvard University Press.


SEDLAC (2014). Socio-Economic Database for Latin America and the Caribbean. Available: <http://sedlac.econo.unlp.edu.ar/eng/statistics.php>. Last accessed 20.03.2015.

Stiglitz, J.E. (2011). Of the 1%, by the 1%, for the 1%. Available at: <http://www.vanityfair.com/society/features/2011/05/top-one-percent-201105>. Last accessed on 04.08.2014.

Tello, C. (2010). *Sobre la Desigualdad en México*. México DF: Universidad Nacional Autónoma de México.

Figures and Tables


Figure 1: Inequality in Mexico 1950–2012 (Development of the Palma Ratio and the Gini Coefficient)




Source: Constructed with data from Sedlac (2014) and Tello (2010).

Figure 2: Income Distribution in 116 countries, by Population Share (2008)

Panel A



Panel B




Panel A: Countries ordered according to the size of their bottom 40 per cent’s income share: the tails define the inequality level, while the middle remains ‘stable’ across countries.

Panel B: Countries ordered according to the size of their top 5 per cent’s income share: the diversity of the top ventile across countries contrasts with the homogeneity of the 19th ventile.

Source: Constructed with data from Milanovic (2014).


Figure 3: The Top of the Income Distribution for 116 Countries(2008)



Note: Income shares held by the 10th decile, and the 19th and 20th ventiles individually.

Source: Constructed with data from Milanovic (2014).


Figure 4: The Palma and the 5/40 ratios (116 countries, 2008)



Note: For the sake of readability, South Africa, with a Palma of 9.9 and a 5/40 ratio of just over 7, is excluded from the chart.


Source: Constructed with data from Milanovic (2014).

Figure 5: 41 Countries Ranked According to the Income Share of their Top 1%



Source: Constructed with data from LIS (2014).


Figure 6: Comparing the Original Palma with the 5/40 and 1/40 ratios (41 countries, latest year)



Note: For the sake of readability, South Africa’s Palma and 5/40 ratios, which exceed 4 by far, are excluded from the chart.


Source: Constructed with data from LIS (2014).

Figure 7: Shares Held by Different Income Groups in Colombia, Pakistan, Philippines, and Thailand (1988–2008)



Source: Constructed with data from Milanovic (2014).

Figure 8: Income Share Held by the Top 1 Per Cent for 25 countries (1990–2010)



Source: Constructed with data from LIS (2014).

Table 1: Measures of centrality and spread for different income groups in 116 countries, latest year

	Hmean	Median	Average	St Dev	C of V
D10	28.546	28.444	29.881	6.811	0.228
v19	10.179	10.217	10.328	1.282	0.124
v20	18.197	18.237	19.553	5.682	0.291
D1–4	16.535	18.654	17.890	4.348	0.243
D5–9	52.042	53.014	52.229	2.995	0.057
M55	62.466	63.072	62.557	2.342	0.037

Note: Hmean = harmonious mean; St Dev = standard deviation; C of V = coefficient of variation; M55 = augmented middle 55 Per Cent.

Source: Constructed with data from Milanovic (2014).

Table 2: Income Distribution in Mexico and Denmark, differences in before and after taxes Gini coefficients,1968–2010

Year	Mexico	Denmark
1968	0.030	0.310
1970	0.029	0.297
1975	0.030	0.250
1980	0.022	0.177
1985	0.014	0.156
1990	0.012	0.176
1995	0.016	0.215
2000	0.016	0.198
2005	0.022	0.198
2010	0.023	0.214

Source: Adapted from Krozer and Moreno-Brid (2014).

Table 3: Measures of centrality and spread for different income groups in 25 countries, around 1990–2010

	Hmean		Median		Average		St Dev		C of V		Change
	1990	2010	1990	2010	1990	2010	1990	2010	1990	2010	1990/2010
D10	22.2	23.3	22.3	23.1	22.7	23.6	4.2	3.0	0.187	0.129	1.11
v19	8.8	9.0	8.9	9.0	8.9	9.1	1.0	0.8	0.108	0.084	0.23
v20	13.3	14.2	13.4	13.8	13.8	14.5	3.4	2.3	0.242	0.161	0.90
P100	4.1	4.6	4.0	4.6	4.5	4.8	1.8	1.1	0.395	0.229	0.49
D1– 4	21.5	20.5	22.4	21.4	22.0	20.9	3.2	2.9	0.147	0.138	-1.00
D5– 9	55.2	55.3	55.3	55.2	55.2	55.3	1.9	1.1	0.034	0.019	0.15
M55	64.1	64.4	63.9	64.2	64.1	64.4	1.6	1.3	0.025	0.020	0.32

Note: Hmean = harmonious mean; St Dev = standard deviation; C of V = coefficient of variation; The numbers provided refer to the year closest to 1990 and 2010 for which data is available.

Source: Constructed with data from LIS (2014).