Department of Economic & Social Affairs

DESA Working Paper No. 143
ST/ESA/2015/DWP/143

July 2015

Inequality and the Tails: The Palma Proposition and Ratio Revisited

*Alex Cobham, Luke Schlogl and Andy Sumner**

ABSTRACT

This paper revisits the earlier assessments of the Palma Proposition and the 'Palma Ratio'. The former is a proposition that currently changes in income or consumption inequality are (almost) exclusively due to changes in the share of the richest 10 per cent and poorest 40 per cent because the 'middle' group between the richest and poorest always capture approximately 50 per cent of gross national income (GNI). The latter is a measure of income or consumption concentration based on the above-mentioned proposition and calculated as the GNI capture of the richest 10 per cent divided by that of the poorest 40 per cent. In this paper we do the following: note the use already being made of the Palma Ratio; consider the issue of hidden (or partially hidden) inequality and how the Palma may be useful in bringing this to light in the parts of the distribution that we are likely to be more interested in (the richest and the poorest); revisit the empirical basis of the Palma Proposition (the relative stability of the 'middle') with a new and expanded dataset across and within developing and developed countries. We find the data reaffirms the Palma Proposition and that the proposition is getting stronger over time. We also discuss the theoretical and empirical questions and implications arising from the Palma Proposition as areas for future exploration.

JEL Classification: D63

Keywords: Inequality; Gini; Palma Ratio

* Thank you to a reviewer for comments on an earlier draft.
All correspondence to King's College, London: andrew.sumner@kcl.ac.uk

CONTENTS

Typesetter: *Nancy Settecasi*

UNITED NATIONS
Department of Economic and Social Affairs
UN Secretariat, 405 East 42nd Street
New York, N.Y. 10017, USA
e-mail: undesa@un.org
http://www.un.org/en/development/desa/papers/

Inequality and the Tails: The Palma Proposition and Ratio Revisited

1 Introduction

This paper revisits the earlier assessments of the Palma Proposition and the 'Palma Ratio' in Cobham and Sumner (2013a; 2013b). In this paper we do the following: note the use already being made of the Palma Ratio; consider the issue of *hidden* inequality and how the Palma may be useful in bringing this to light in the parts of the distribution that we are likely to be more interested in (the richest and the poorest); revisit the empirical basis of the relative stability of the 'middle' with a new and expanded dataset across and also within developing and developed countries.

The Gini coefficient remains the dominant measure of income or consumption/expenditure inequality. However, interest has emerged in tracking top incomes (and in some cases adjusting the Gini with these), most notably in the work of Emmanuel Saez, Thomas Piketty and Tony Atkinson and others at the *World Top Incomes Project* who have sought to track, through tax records, the incomes of the richest 10 per cent, 5 per cent, 1 per cent, 0.1 per cent and 0.01 per cent and other fractiles in 30, largely but not entirely, Organisation for Economic Co-operation and Development (OECD) countries with available tax data (see Alvaredo et al., 2014).

An alternative approach to tracking changes in inequality related to the top incomes in society is that of José Gabriel Palma (2006, 2011, 2013, 2014a, 2014b). Palma has argued empirically that changes in inequality are currently determined by the richest 10 per cent and the poorest 40 per cent because the population in-between (deciles 5–9) hold a relatively stable half of gross national income (GNI) irrespective of country and time (relative that is to the stability of the shares to the poorest 40 per cent or the top 10 per cent).

Based on this proposition, the 'Palma Ratio' of income or consumption concentration is a measure of the capture of total income or consumption of the richest decile over the capture of the poorest 40 per cent. To take a focus on the top incomes, as the Palma Ratio does and Piketty and others do, inevitably raises questions about the quality of the data at the top end of the distribution, where under-sampling (and under-reporting) of richer households is a serious problem. Piketty and others address this issue by using tax records rather than survey data as we use below. The Palma Ratio could be based or adjusted for top incomes data as more become available; it can also be adjusted for untaxed income. The more commonly used and technically (axiomatically) stronger, Gini measure of inequality is, however, as dependent on survey data, given tax data only exist for 30 countries in the *Top Incomes Project* database.

Policy-related interest (meaning citation and usage) in the 'Palma Ratio' has grown over the last one to two years, driven, in part, by proposals to include a 'Palma target' in the UN's post-2015 framework for global development such as that by Joseph Stiglitz (see Doyle and Stiglitz, 2014).[1] Data for the Palma Ratio is now listed and updated as standard measure of inequality in the OECD Income Distribution database (see Cingano, 2014 and OECD, 2014) and the UNDP annual Human Development Report (See UNDP, 2014), as well as by some national statistical offices, e.g. the UK (ONS, 2015).[2] Further, interest in the Palma Ratio is evident among NGOs

1 Doyle and Stiglitz propose a Palma Ratio of 1 by the year 2030. In contrast, Engberg-Pedersen (2013) suggested a more contextualised approach, in which countries would aim to halve the gap between their starting point and a Palma of 1 by 2030.

2 One could go as far as to say that the Palma Ratio might be implicit in the World Bank's goal of 'shared prosperity' which although focuses on the growth of incomes of the bottom 40 per cent, implicitly will focus on the richest 10 per cent if the Palma proposition holds that deciles 5–9 capture 50 per cent of GNI.

and international agencies alike (see for illustration, EC, 2014; OECD, 2014; Oxfam, 2014; UNDESA, 2013).

In this paper we find that the latest data not only supports the Palma Proposition but that the proposition is getting stronger over time. We also discuss the theoretical and empirical questions and implications arising from the Palma Proposition for future exploration.

The paper is structured as follows: Section 2 discusses the notion of hidden inequality and how the Palma may be more useful than the Gini in some aspects. In Section 3 we revisit and expand a set of tests on the Palma Proposition that is the basis of the Palma Ratio. In Section 4 we discuss theoretical and empirical questions arising. Section 5 concludes. An annex provides further conceptual discussions drawing from Cobham and Sumner (2013a; 2013b). A second annex provides further tests.

2 Three Dimensions of Hidden Inequality

Various conceptual issues and a set of policy-based axioms for the Palma Ratio were discussed in Cobham and Sumner (2013a; 2013b) (and are included as an annex in this paper). In this section we take as the point of departure methodological issues related to the area of hidden inequality. These arise in three main ways: due to data constraints; due to the hiding of particular behaviour; and due to potentially misleading choice of measures. The first two of these dimensions of 'hidden inequality' are factors with regard to all measures of inequality. However, the third dimension of hidden inequality forms the basis for the use of the Palma Ratio.

The quality and availability of data on income and consumption inequality remains a constraint. As alluded to at the outset of this paper, with regard to top incomes, the weaknesses of income distribution data are well known. Carr-Hill (2013) identifies systematic under-sampling and exclusion by design of marginalised, likely poorer populations from both survey and census data. At the other end of the distribution, there is clear evidence of non-response from high-income groups. Korinek, Mistiaen and Ravallion (2005) apply a model to the Current Population Survey of the United States, and find that the probability of compliance falls monotonically as income increases; and that this results in appreciable understatements of both mean income and of inequality: the upward corrections of the Gini during the sample years 1998–2004 range from 3.39 to 5.74 percentage points (raising the Gini from around 0.45 to 0.49–0.50).[3]

Anand and Segal (2014) extend the same intuition to global distributions using the *World Top Income Database* (and global inequality between individuals, on the basis of per capita income in their household, using 2005 PPP$ for international comparison). Starting with Milanovic's (2012) data set of five-yearly observations, and assuming that it represents the bottom 99 per cent, they scale accordingly and add the recorded top 1 per cent income (which also increases mean income). The database covers 104 country-years of the 537 in the sample, and so the missing top percentile income shares are imputed on the basis of regression. The results are broadly consistent with national findings, although the major shift in top income shares when top income data are included is not well captured in the Gini.[4]

The second hidden area of inequality relates to income and wealth which is simply undeclared – so that it is absent from tax administrative data as much as from survey responses. Globalisation has opened up opportunities in this area, as cross-border economic and financial ownership has become common but national mechanisms of regulation have failed to keep pace. Major impediments include banking secrecy, the use of anonymous companies, trusts and

3 Atkinson, Piketty and Saez (2011) use top income data to reconsider the same issue, and find similar results: for 2006, the Gini is 4.9 percentage points higher, at 0.519 as opposed to the original calculation of 0.470.

4 The biggest change occurs in 2005 and is four percentage points – whereas Theil's T, for example, shows an equivalent change of 22 percentage points.

foundations, other aspects of poor corporate transparency and the relative weakness of international information exchange. Zucman's (2013) estimate is that undeclared 'offshore' financial assets may amount to nearly 10 per cent of world GDP. The scale of estimates suggests that adjustment could potentially have large impacts on national income means and inequality: for example, the African Development Bank and Global Financial Integrity (2013) estimate that 18 different African countries saw average annual illicit financial outflows in excess of 10 per cent of GDP between 1980 and 2009. Some approaches (for example, Zucman, 2013) estimate undeclared wealth and then construct estimates of the resulting undeclared income. Others estimate undeclared incomes, in the form of illicit financial flows, and use these to construct undeclared wealth estimates (for example, Boyce and Ndikumana, 2013).

Such illicit financial flows can include not only tax abuse (individual and corporate tax evasion and – possibly – some avoidance), but also laundering of the proceeds of crime, the theft of state assets and the bribery of public officials, and some forms of market abuse (for example, use of anonymous ownership to circumvent antitrust regulation or to hide political conflicts of interest).

The two previous dimensions of 'hidden inequality', outlined above, affect all measures of inequality. However, there is a third dimension to hidden inequality which forms the basis for the use of the Palma Ratio: Inequality has also been hidden by the choice of measure itself. The Atkinson (1970) critique highlights the particular weaknesses of the Gini coefficient as a measure of inequality, noting that the Gini 'attaches more weight to transfers affecting middle income classes' (pp. 256–57). In short, the Gini is relatively oversensitive to changes in the middle (relative to changes at the extremes of the distribution; and also relatively insensitive at higher levels of inequality (see Cobham and Sumner, 2013a; 2013b) Does it matter that this is not explicit? What does one care about – the distribution in the middle or at the extremes? And when does one care – at low or high levels of inequality? And what if changes to the middle tend to be limited in practice, as we show in the following section? That would mean that using the Gini would be to choose a measure of inequality that is most sensitive to changes that are less common, in a part of the distribution that we might be less concerned about, while being under sensitive to the part of the distribution where change is more likely, and which we might be more concerned about. On top of this, that the measure in question is insensitive at higher inequality levels, and does not make any of these normative judgements explicit.

These questions are far from academic – in addition to the data constraints and problems of hidden behaviour discussed above, the choice of inequality measure can substantially change the view that emerges of a particular situation. We can illustrate the above points with UK inequality data. Figure 1 compares the Gini and Palma for household incomes before housing costs for 1961-2012/13, indexed to the start of the series to assist comparison. The most well-known feature of UK inequality – the major increase from the late 1970s to late 1990s – is confirmed by both measures. The subsequent divergence, however, supports quite different views. By the Gini, inequality from in late 1990s remained around the same level to 2010 more or less. By the Palma, however,

Figure 1
UK Income Inequality, 1961-2011/12 by Gini and Palma Ratio

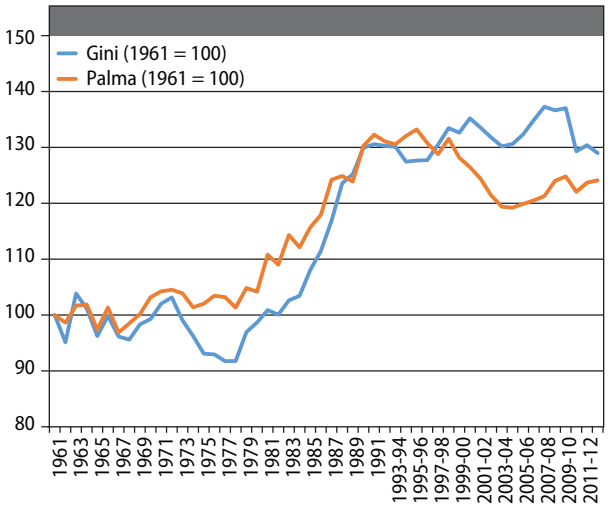

Source: Authors' estimates based on data from IFS (2015).

the same period registered a reduction in income inequality which is in contrast to the earlier period. There is also a period in the early 1970s where Palma and Gini diverge (Gini declining, Palma steady).

Whether the Gini story or the Palma is 'right' is a normative question. What the example illustrates how the choice of measure shows different patterns.

3 Income and Consumption Inequality and the Tails

a. The Palma Ratio

The Palma Ratio is a particular specification within a family of inequality measures known as inter-decile ratios, such as the GNI share of the poorest 20 per cent of the population over the richest 20 per cent, (or its inverse). However, rather than the GNI share of the bottom 20 per cent over the top 20 per cent, the Palma Ratio is the ratio of national income shares of the top 10 per cent of households to the bottom 40 per cent on the basis of José Gabriel Palma's (2006, 2011, 2013, 2014a, 2014b) empirical observation of the cross-country stability of the 'middle' share of income so that distribution is largely a question of the 'tails' of the distribution (the poorest and the richest).

The 'Palma Proposition' (originally outlined in Palma, 2006) is as follows: changes in income inequality are exclusively due to changes in the share of the richest (D10) and poorest (D1–4), leaving unchanged the income share of the 'middle', meaning the 5th to the 9th decile group (D5–9). On that basis, one could argue that half of the world's population (the middle and upper-middle groups) have currently acquired strong 'property rights' as Palma puts it, over half of their respective national incomes, while there may be more flexibility over the distribution of the other half of this income, between the 'rich' (the top 10 per cent of the population) and the 'poor' (the bottom 40 per cent of the population). Palma established an argument that there are two opposite forces at work on the national distribution of each country: one force is 'centrifugal', leading to a divergence in the shares of the top 10 per cent and bottom 40 per cent, the other force is 'centripetal', leading to a convergence in the income share appropriated by the 'middle' 50 per cent (deciles 5 to 9). 'Centripetal forces', or growing uniformity implies (by definition) that there is change, and change with a direction. In a recent note, Hazledine (2014) provides a critique of this 'Palma Proposition' which Palma (2014a) replies to. Hazledine pursues the question of the rigidity of the middle 50 per cent income share. Palma replies, reiterating earlier papers that it is the *relative* stability of the middle 50 per cent. Nobody has argued that the middle share is fixed, which would be trivially false; the argument is about the *relative* stability of this group's income share (compared to the richest or poorest). Annex II of this paper provides further tests and compares the stability of the 'middle' across the coefficient of variation (the ratio of the standard deviation to the mean) with the standard deviation and other metrics of statistical dispersion, specifically the interquartile range and median absolute deviation. We find that all of these show that middle shares have a high degree of cross-country homogeneity. For this reason, in the text of this paper we use the coefficient of variation.

A further critique is contained in a blog by Milanovic (2015) which focuses the discussion of the Palma Proposition in Krozer (2015). The critique is that the Palma Ratio as a general measure of inequality is an empirical regularity that may not hold in the future. Milanovic also reiterates two points in respect of the typical technical axioms for inequality measures, which is that the Palma is insensitive to transfers within any of the three 'chunks' (the top 10%, middle 50% or bottom 40%) and the decomposition properties. These latter points have been discussed in Cobham and Sumner (2013) and are contained in the annex of this paper for reference.

First, one should ask what would an appropriate test of the constancy of shares look like? A simple comparison of the variation of shares is one way to determine if the middle is still homogenous and if the outer shares still show higher variation. Cobham & Sumner (2013a; 2013b) do this. In addition to that, we can examine the frequency and distribution of changes of shares and consider measures of the

central tendency and variability of change. Data below show that stability over time regarding all of Palma's decile groups. In comparison, though, D5-9 has, on average, clearly shown least change *over time* namely only 0.27 percentage points on average (and it also shows the lowest variability of change across countries). This is to be expected or else it would have been a coincidence that the cross-country variation of D5-9 income shares had remained so homogenous. Indeed, this pattern of temporal stability is already suggested by the initial cross-country homogeneity which implies that countries with very different characteristics and at very different levels of development exhibit a similar D5-9 share. Where should (sudden) change in the middle capture thus come from? Bearing in mind the strong cross-country homogeneity it is therefore not unreasonable to expect a (high) degree of inter-temporal stability of shares. The rate of change of D5-9 shares across time centres around zero with a median that is slightly negative around -0.23 percentage points of change (see figure 5 below).

The Palma implicitly contains information about the share of the middle and thus is not blind to changes of the *aggregate* D5-9 share. However, the question is one of sensitivity towards changes happening *in the middle of the income distribution*. In contrast to the Gini the Palma lumps together the incomes of half of the population in the 'middle' of the income distribution. The Palma thus *is* 'blind' towards intra-middle variation. This shortcoming might be offset by other desirable characteristics it has as a measure of income concentration. A normative justification for the use of 'the Palma' was advanced by Cobham and Sumner (2013a; 2013b). So the Palma might be the right measure for today; but what if the empirical regularity were to cease to hold in the future? It is still the case that it is the top decile capture that matters. The correlations between the Palma and ratios of the top decile to other parts of the distribution confirm this. The Palma Ratio which should be understood as measure of income concentration, is not sensitive at all to the selection of the bottom deciles. Table 3 below shows the correlation of the Palma Ratio with ratios of the richest decile over different decile configurations. It becomes evident that, with the exception of D8-9 and D9, it makes little difference which decile combination is chosen for the denominator of the Palma. Debating which deciles to include or exclude is thus secondary, Even if the future does hold a sufficiently dramatic change in the relative stability of the deciles 5-9 what would we be left with? The Palma as a measure of income or consumption concentration, which: remains meaningful; pays sufficient attention to a part of the distribution that we may care about; and is explicit about doing so. The Gini is by construction oversensitive to the middle, and less sensitive to the tails. As such, it is an inequality measure which: remains meaningful; pays insufficient attention to a part of the distribution that we *do* care about; and is not explicit about doing so. We would suggest that most people using the Gini do not realise that it is less sensitive to the tails; nor that it becomes increasingly less sensitive at higher levels of inequality. As such, use of the Gini can hide the true extent of inequality – inadvertently or otherwise. We should recognise there are weaknesses to *any* single measure of inequality. As Atkinson, wrote in 1970, all measures reflect a subjective view – the difference is whether this is made explicit.

In the next section we revisit earlier tests in Cobham and Sumner (2013a, 2013b) using the latest and largest dataset available which is the World Bank (2015) and confirm (as in earlier tests), Palma's proposition of the *relative* stability of the middle 50 per cent income share not only across countries but – additionally – within countries over time. We find that the Palma Proposition if anything is getting stronger.

b. Empirical tests

First, we test the robustness of Palma's 'middle capture' across countries in 1990 and 2012 and a combined sample (Table 1); and for within-country variance (Table 2). To assess the Palma Proposition we use decile data on income/consumption distribution from the World Bank (2015) PovcalNet data set (March 2015). We take data for the nearest dates to 1990 and 2012 using the respective windows of 1985–1995 and 2005–2012. The combined data set (1990 and 2012) includes 141 countries in total. We

consider income and consumption shares in both periods separately and combined. Because no means of adjustment (income vs consumption) is readily acceptable we do not adjust surveys, but consider country changes only by looking at surveys of the same type.

i. The relative stability of the capture of the Palma 'middle'

We confirm that Palma's finding of the relative stability of the capture of GNI of the 'middle 50 per cent' (deciles 5 to 9) holds in both 1990 and 2012 (See Table 1). The 'middle' share mean is 0.52 across the observations. Although there is a minimum of 30.67 and a maximum of 57.27 for the 'middle' capture the coefficient of variation is just 0.08. If one considers all observations, almost 90 per cent of all observations for the 'middle' are above 45.00 or below 56.00. The coefficient of variation of the 'middle' compares to 0.27 in the poorest four deciles (mean 0.31) and 0.26 in the richest decile (mean 0.18). In short, the relative variance of the 'middle' is substantially lower than the richest decile or poorest four deciles. The 'middle' does capture half of GNI on average and the richest 10 per cent capture, on average, three times their population share while the poorest 40 per cent population capture half of their population share. The values for mean, minimum and maximum do not differ very much if one exclusively takes the consumption shares or the income shares or focuses solely on 1990 or 2012. The only notable difference between 1990 and 2012 is across each of the three population groups (the richest decile, poorest four deciles and middle five deciles): the coefficient of variation has fallen over time in each group suggesting some convergence in the capture of each share. In short, the stability of shares has increased over time. The 'middle' share varies consistently less across countries than do the shares of the top 10 per cent and bottom 40 per cent; all three are more stable across countries in 2012 than in 1990, but the 'middle' has a coefficient of variation which is consistently a third of that of the top 10 per cent or the poorest four deciles. When we consider the coefficients of variation according to the type of survey (income versus consumption), the stability of the middle is also confirmed.

In Table 2 we consider temporal stability within countries. We calculate the grand mean (that is, the sum of all observed values divided by the total number of observations). The grand mean gives equal weight to every observation (which is usually desirable but might perhaps not be in our case) and the mean of means (that is, a mean over all countries' mean values). This gives equal weight to every country irrespective of whether it has data for one year or for 20. Again we find the Palma Proposition holds – the relative stability of the Palma 'middle' is much lower (the coefficient of variation) than that of the poorest 40 per cent or richest 10 per cent.

What is the reason for the close to 50% share of D5-9? It has to do with the homogenous geometry of income distributions across countries rather than with the arithmetic of decile groupings. The capture of shares of deciles 1 to 10 is strongly skewed to the right in most countries and the largest percentage point variation can be observed in the share of the top decile (Figure 2). It is visible with the naked eye that the share of the upper middle deciles, in particular, show very little dispersion (with one outlier at the bottom belonging to the Seychelles). However, it is true that

Figure 2
Consumption shares of 1st to 10th decile across countries (2012 or nearest data available)

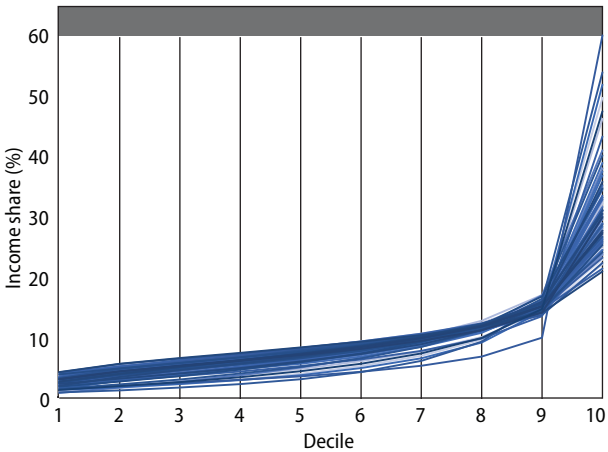

Source: Authors' estimates based on data from World Bank (2015).

the bottom decile captures show more similarity with the middle deciles than with the top deciles.

If we take this as stylized fact and consider the income share of D5-9 across different (hypothesized) exponential income distribution scenarios, we find that this share first increases after complete equality, while the share of D1-4 falls (Figure 3). D5-9 reaches a turning point around 60% capture and then declines smoothly converging towards zero while the top share converges against 100%. The turning point of the D5-9 share is located at 1.2 (the point where the next decile gets 1.2 times the income of the previous decile). Due to the turning of the share curve, there are ample opportunities for the D5-9 share to take values around 50%. The exponential structure of most income distributions might be the reason why we observe a maximum of 55.83% of D5-9 capture in current WDI consumption data and homogeneity in the middle. Captures beyond 60% might require rare or unrealistic income distributions (such as logarithmic distributions or distributions with discrete jumps). As noted earlier, the rate of change of D5-9 shares across time centres around zero with a median that is slightly negative around -0.04 percentage points of change (figure 4).

Figure 3
Share of D5-9 across a range of exponential distribution scenarios

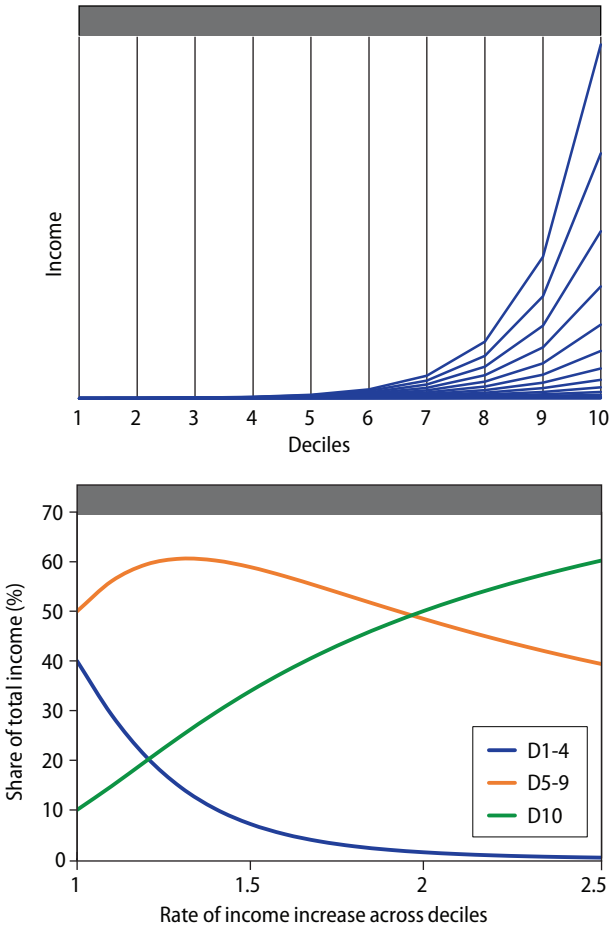

Source: Authors' calculations based on hypothetical distribution scenarios.

Figure 4
Frequency distribution of percentage point changes of D5-9 income or consumption shares (1990-2010 or nearest data available)

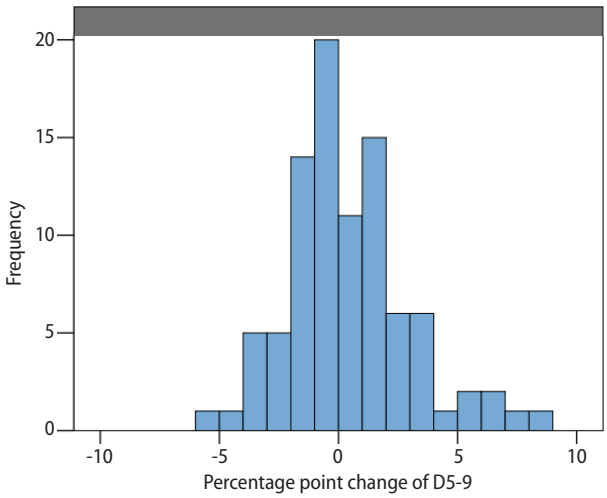

Source: Authors' estimates based on data from World Bank (2015).

Cobham & Sumner (2013) show the Palma is quite closely correlated with the income share of the middle 50 per cent and that, in general, higher Palma Ratios imply a squeezing of the share of the middle 50%; so in practice the Palma will tend to reflect income concentration here too, even though it is not directly captured in the ratio. This is confirmed by the slopes of the regression lines in Figure 5. The rates of change of D5-9 against D1-4 and D10, respectively, are inverted, suggesting that the 'middle' increased their shares in situations where the poor

Figure 5
Changes of income or consumption decile shares over time of D1-4 vs D5-9 and D10 vs D5-9 (1990-2012 or nearest data available)

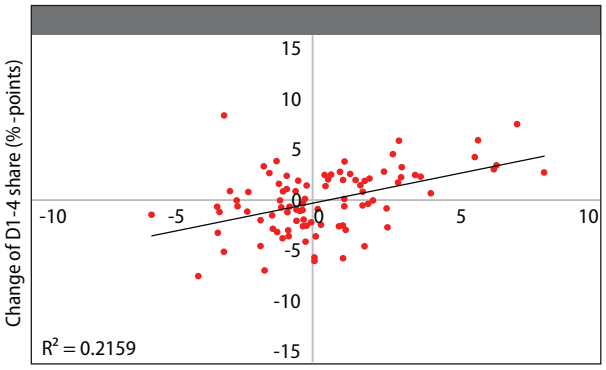

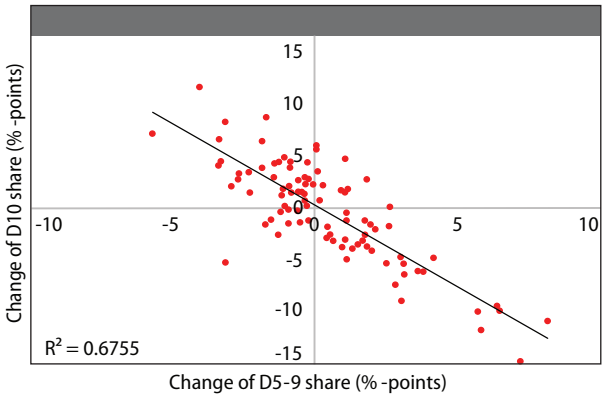

Source: Authors' estimates based on data from World Bank (2015).

increased their shares as well and that the middle's shares decreased in situations where the shares of the richest increased.

A final point to reiterate is that the Palma *is* 'blind' towards intra-middle variation. However, this shortcoming might be offset by other desirable characteristics it has as a measure of income concentration. We know that *by construction* the Gini is over-sensitive to the middle; but in practice it is equally *in*sensitive to the middle as is the Palma; so the implication would seem to be that the Palma Proposition holds sufficiently strongly to overcome the Gini's bias (possibly exacerbated by weaknesses in constructing Gini series from limited quantile data). That leaves a choice between a measure which by design is oversensitive to the 'wrong' bit of the distribution, but in practice tells us nothing about it; and a measure which by design and practice, deliberately tells us nothing about it. If you want to know about the middle, the Gini seems to be little good to you – but may fool you that it is.

In sum, we can find no evidence to undermine the Palma Proposition. There are, however, a number of theoretical and empirical questions arising from the Palma Proposition and Palma Ratio which we discuss next.

4 Theoretical and Empirical Questions

There are a number of theoretical and empirical questions arising relating to the Palma Proposition and Palma Ratio. We discuss three areas that we believe to be particularly pertinent. First, there are questions relating to the conflation of 'middle' groups across all developing countries (and the use of decile income or consumption to discuss 'class' which is a social identity). Second, there are questions relating to intra-'middle' (and intra-poorest 40 per cent or intra-richest 10 per cent) shifts. Third, there are questions relating to the theoretical mechanisms underlying the Palma Proposition.

First, the conflation of the 'middle' across all countries. The Palma Proposition refers to deciles 5 to decile 9 as the 'middle' (with deciles 5 and 6 the lower middle and deciles 7-9, the 'upper middle'). Not only is that not the literal middle, if one took the UK, Brazil and Burundi, the 'middle' (i.e. deciles 5 to 9) would be quite different levels of income or consumption with much of the middle in the latter living in absolute poverty. One cannot, of course, conflate social identity and income/consumption data in more than the most general sense because in some countries the 'poor' will be in the middle deciles. Palma (2011, p. 102) is clearly aware of this issue himself. He suggests that, in light of the observation that the share of GNI of those people in deciles D5–D9 is generally half of national income, the 'middle classes' should be renamed the 'median classes': 'Basically, it seems that a schoolteacher, a

junior or mid-level civil servant, a young professional (other than economics graduates working in financial markets), a skilled worker, middle-manager or a taxi driver who owns his or her own car, all tend to earn the same income across the world – as long as their incomes are normalised by the income per capita of the respective country.'

We would go further and suggest that the main point is that changes in inequality are determined by what happens to the richest decile and poorest four declines and thus if the middle represents different consumptions or income ranges in different countries, as it will, that does not undermine the basic proposition or use of the Palma Ratio.

Second, of more significance to measuring inequality are intra-'middle' shifts. The major drawback of the Palma Ratio is, of course, that it is not a measure of the whole distribution but excludes deciles 5–9 in the middle (and movements within the poorest 40 per cent or top 10 per cent). Relatedly, the Palma does not meet the standard axioms for inequality measures: in particular, Pigou-Dalton transfer sensitivity is violated as noted in Cobham and Sumner (2013a; 2013b), where a transfer from, say, the 89th centile to the 41st fails to result in a lower Palma. Palma (2011) too has pointed towards a potential difference within the 'middle' between the GNI capture of the lower middle, or D5–D6 versus the upper middle, D7–D9 (and a very large difference between D9 versus D10 capture of GNI). In response, one could view the Palma Ratio not as a measure of inequality per se but as a measure of income or consumption concentration. The fact that the Palma Ratio does not allow for capturing the intra-group churning, while the Gini does (even if it is more sensitive to changes in the middle of the distribution and less sensitive at the extremes), is important if the objective is a measure of the entire distribution. If, however, the objective of using the Palma Ratio is not to consider the entire distribution and make statements on inequality in entirety, but rather focus on the extremes under the explicit and normative judgement that it is changes in consumption or income concentration not inequality per se and the tails in particular that matter, then the Palma Ratio would be a useful measure to take. Also, if the share of D5-D9 across countries today is fairly homogenous, the Palma Ratio would suffice for international comparisons of inequality, as almost all diversity of inequality across countries takes place among D10 and D1-D4.

Third, one could ask what the mechanisms are by which the 'median class' is able to control a fairly stable or rather, a fairly homogenous share of the national income. Are such mechanisms the same in countries with very different characteristics? Or, if they are different, what leads to a quasi-universal similarity in their income share? Palma (2014b) himself posits that there is a 'sub-optimal equilibrium' with regard to a specific group of Latin American middle income countries where inequality is high but the decide 5 to decile 9 still get their share, which is as follows: the situation is more stable than one would expect even in a democracy because the rich do well, the 'middle' have access to cheap services (for example, domestic maids) and expanding (service sector) employment, albeit poorly remunerated, helps the poorest. As a result, in Latin America, high inequality exists with low growth and low unemployment. This contrast with, say, South Africa, where inequality is such that D5-D9 does not get enough even to afford cheap services, leading to high inequality with high unemployment.) Palma discusses the high inequality, low unemployment, low growth, equilibrium thus:

It keeps the rich blissful (huge rewards with few market 'compulsions'); it allows the middle and upper middle groups to have access to a particularity large variety of cheap services; and it does at least provide high levels of employment for the bottom 40 per cent... jobs may be precarious, mostly at minimum wages… and in activities with little or no potential for long-term productivity growth, but at least they are jobs and there are plenty of them (pp. 28–29).

What requires further theory building would be how the 50 per cent of GNI came to be captured in the first place. That would require some major undertaking of the history of income and consumption inequality in a range of developing countries to identify

commonalities. Nel (2012, pp. 24–29) proposes such a historical framework based on asset concentration, the mode of incorporation into the world economy, economic modernisation and governance, among other factors. If there is one area for taking forward the discussion of the Palma Proposition, this would seem to be an important one.

In sum, one could argue that the Palma Proposition has two sides: One is that if it is true that today D5-D9 gets – and can defend – their 50% of income, this has important implications if one wants to reduce inequality. The other is how did this situation arise? What are the forces that led to this convergence in the middle and upper-middle across countries? Even if we do not know much about the dynamic that led to the latter, the knowledge of former issue in itself could be relevant for the targeting of redistributive policies, such as fiscal policy.

5 Conclusions

In this paper we have revisited the Palma Proposition and relatedly, the Palma Ratio as a measure of consumption or income concentration. In conclusion we can say (again) the Palma Proposition holds in the sense of the relative stability of the 'middle' or 'median' group between the poorest 40 per cent and richest 10 per cent and thus the Palma Ratio of income or consumption concentration has two main attractions.

First, the Palma focuses attention on the aspect of inequality that has to a number of researchers and policy debates been identified as important: that is, the relationship between the tails or the richest and the poorest. If that is the explicit normative choice the Palma not only avoids the Gini's oversensitivity to the middle, but also the Gini's relative insensitivity to changes at the top or the bottom of the distribution (where one might expect policymakers to be *more* concerned rather than less). If one wanted a measure of the entire distribution then the Gini would be more appropriate as that is what the Gini does and the Palma Ratio does not.

Second, one might say that the primary strength of the Palma Ratio is its simplicity for use in policy debate: a Gini coefficient of 0.5 implies serious inequality but yields no intuitive statement for a non-technical audience. In contrast, the equivalent Palma of 5.0 can be directly translated into the statement that the richest 10 per cent earn five times the income of the poorest 40 per cent of a country. That simplicity we would argue makes the Palma Ratio useful in policy debates, at least as useful as the Gini in tracking income/consumption and its concentration. One could argue that on top of simplicity, the Palma Ratio also helps to focus the policy debate on inequality on what can be done effectively. As mentioned above, if D5-D9 already succeeds in capturing 50 per cent but they are somehow stuck at that limit, to do something significant about inequality one has to do something at the tails; the Gini is not only obscured to the non-technical audience: it has also obscured so far this fact to the technical audience). Perhaps the most important messages from the rise of the Palma Ratio in policy discussions are the need to be aware of implicit normative aspects embedded in particular measures; and the risks of allowing one measure to dominate. In particular, the Gini, as the currently dominant measure but with an insufficiently appreciated bias, contributes to a wider problem of 'hidden' inequality.

References

African Development Bank/Global Financial Integrity (2013), *Illicit Financial Flows and the Problem of Net Resource Transfers from Africa: 1980-2009*, Tunis: African Development Bank.

Alvaredo, F., Atkinson, A., Piketty, T. and Saez, E. (2014) The World Top Incomes Database [online]. Available from : http://topincomes.g-mond.parisschoolofeconomics.eu [Accessed 14 May 2015].

Anand, S. and Segal, P. (2014) *The Global Distribution of Income*. Oxford Department of Economics Discussion Paper 714. Oxford: Oxford University.

Atkinson, A. (1973) 'On the Measurement of Inequality', reprinted from *Journal of Economic Theory* (1970), with non-mathematical summary, pp. 46–68, in A. Atkinson (ed.), *Wealth, Income and Inequality*. Harmondsworth (UK): Penguin Education.

Atkinson, T. (1970) 'On the Measurement of Inequality', *Journal of Economic Theory*, 2, pp. 244–63.

Atkinson, A., Piketty, T. and Saez, E. (2011) 'Top Incomes in the Long Run of History', *Journal of Economic Literature*, 49 (1), pp. 3–71.

Boyce, J. & L. Ndikumana, 2012, 'Capital flight from Sub-Saharan African countries: Updated estimates, 1970 – 2010', *Political Economy Research Institute (PERI) Research Report*. University of Massachusetts Amherst.

Carr-Hill, R. (2013) 'Missing Millions and Measuring Development Progress', *World Development,* 46, pp. 30–44.

Cingano, F. (2014) *Trends in Income Inequality and its Impact on Economic Growth*. OECD Social, Employment and Migration Working Papers 163. Paris: OECD.

Cobham, A. and Sumner, A. (2013a) 'Putting the Gini Back in the Bottle? "The Palma" as a Policy-Relevant Measure of Inequality'. Mimeograph. London: King's College London.

Cobham, A. and Sumner, A. (2013b) *Is It All About the Tails? The Palma Measure of Income Inequality*. CGD Working Paper. Washington DC: CGD.

Cowell, F.A. (2007) *Measuring Inequality* (third edn). Hemel Hempstead: Oxford University Press.

Cowell, F.A. (2000) 'Measurement of Inequality', in A.B. Atkinson and F. Bourguignon (eds), *Handbook of Income Distribution*. Amsterdam: North Holland, pp. 87–166.

Dalton, H. (1920) 'The Measurement of the Inequality of Incomes', *Economic Journal*, 30, pp. 348–61.

Doyle, M. and Stiglitz, J. (2014) 'Eliminating Extreme Inequality: A Sustainable Development Goal, 2015–2030', *Ethics and International Affairs*, 28 (10), pp. 1–7.

European Commission (EC) (2014) *EU Employment and Social Situation Quarterly Review: Towards a Better Measurement of Welfare and Inequalities*. Brussels: EC.

Engberg-Pedersen, L. (2013) *Development Goals Post 2015: Reduce Inequality*. Danish Institute for International Studies Policy Brief. Copenhagen: DIIS.

Duro, J. (2008) 'Cross-Country Inequalities in Welfare and its Decomposition by Sen Factors: The Virtues of the Theil Index', *Applied Economics Letters*, 15 (13), pp.1041–45.

Frosini, B. (2012) 'Approximation and Decomposition of Gini, Pietra–Ricci and Theil inequality measures', *Empirical Economics*, 43 (1), pp. 175–97.

Greselin, F., Pasquazzi, L. and Zitikis, R. (2013) 'Contrasting the Gini and Zenga Indices of Economic Inequality', *Journal of Applied Statistics*, 40 (2), pp. 282–97.

Hazledine, T. (2014) 'Does the Centre Hold? Testing Palma's Proposition (A Comment)', *Development and Change*, 45 (6), pp. 1409–15.

IFS (Institute for Fiscal Studies) (2015) IFS Inequality and Poverty Dataset. Downloaded 9 June 2015 from: www.ifs.org.uk/tools_and_resources/incomes_in_uk

Korinek, A., Mistiaen J. and Ravallion, M. (2005) *Survey Nonresponse and the Distribution of Income.* World Bank Working Paper 3543. Washington DC: World Bank.

Lustig, N., Gray-Molina, G., Higgins, S., Jaramillo, M., Jiménez, W., Paz, V., Pereira, C., Pessino, C., Scott, J. and Yañez, E. (2012) *The Impact of Taxes and Social Spending on Inequality and Poverty in Argentina, Bolivia, Brazil, Mexico and Peru: A Synthesis of Results*. Commitment to Equity Working Paper 3. New Orleans: Tulane University.

Milanovic, B. (2012) 'Global Inequality Recalculated and Updated: The Effect of New PPP Estimates on Global Inequality and 2005 Estimates', *Journal of Economic Inequality*, 10 (1), pp. 1–18.

Milanovic, B. (2015) Why I am Still Not Excited by the Palma. Blog downloaded from: http://glineq.blogspot.co.uk/2015/06/why-i-am-still-not-excited-by-palma.html

Nel, P. (2012) *The Politics of Economic Inequality*. Basingstoke: Palgrave Macmillan.

ONS (2015) The Effects of Taxes and Benefits on Household Income, Financial Year Ending 2014

[online]. Available from: http://www.ons.gov.uk/ons/rel/household-income/the-effects-of-taxes-and-benefits-on-household-income/2013-2014/index.html [Accessed 14 July 2015].

OECD (2014) OECD Income Distribution Database (IDD) [online]. Available from: http://www.oecd.org/social/income-distribution-database.htm [Accessed 14 May 2015].

Oxfam (2014) *Even It Up: Time to End Extreme Inequality*. Oxford: Oxfam.

Palma, G. (2014a) 'Has the Income Share of the Middle and Upper-Middle Been Stable Around the "50/50 Rule", or has it Converged Towards That Level? The "Palma Ratio" Revisited', *Development and Change*, 45 (6), pp. 1416–48.

Palma, G. (2014b) 'Why is Inequality so Unequal across the World? Could it

be that Every Nation Gets the Inequality it Deserves?', in J. Stiglitz and K. Basu

(eds), *Proceedings of IEA Roundtable on Shared Prosperity and Growth*.

Palma, G. (2013) *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 1437. Cambridge: University of Cambridge.

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

Palma, G. (2006) *Globalizing Inequality: 'Centrifugal' and 'Centripetal' Forces at Work*. UN DESA Working Paper 35. New York: UN DESA.

Pigou, A.C. (1912) *The Economics of Welfare*. London: Transaction Publishers.

Schröder, M. (2015) 'Mythical Measures: The Problem of Objective Inequality Measurement in Economics and the Social Sciences', *Groundings* 8: 129–44.

Sen, A. (1973) *On Economic Inequality*. New York: Norton.

Shaheen, F. (2014). *Reducing economic inequality as a Sustainable Development Goal: Measuring up the options for beyond 2015*. London: New Economics Foundatio.

UNDESA (2013) *Inequality Matters: Report on the World Social Situation 2013*. New York: UNDESA.

UNDP (2014) *Human Development Report*. New York: UNDP.

World Bank (2015) *PovcalNet (March)*. Washington DC: World Bank.

Zucman, G., 2013, 'The missing wealth of nations: Are Europe and the U.S. net debtors or net creditors?', *Quarterly Journal of Economics* 128(3), pp.1321-1364.

Table 1
Cross country variance of income and consumption shares, 1990 and 2012 and combined sample

	Richest 10%	Poorest 40%	Middle 50%
Combined sample (1990 and 2012)			
Income and consumption shares (141)			
Mean	0.31	0.18	0.52
Coefficient of variation	0.26	0.27	0.08
Min	17.82	6.74	30.67
Max	59.86	27.70	57.27
Consumption shares only (99)			
Mean	0.32	0.17	0.51
Coefficient of variation	0.24	0.25	0.08
Min	20.03	6.74	30.67
Max	59.86	25.19	56.93
Income shares only (66)			
Mean	0.29	0.18	0.53
Coefficient of variation	0.29	0.29	0.07
Min	17.82	7.10	39.76
Max	50.68	27.70	58.22
1990			
Income and consumption shares (102)			
Mean	0.30	0.18	0.52
Coefficient of variation	0.30	0.31	0.08
Min	17.82	6.92	39.76
Max	51.20	27.70	56.93
Consumption shares only (51)			
Mean	0.34	0.16	0.50
Coefficient of variation	0.23	0.29	0.08
Min	20.03	6.92	39.90
Max	51.20	25.19	56.93
Income shares only (60)			
Mean	0.28	0.19	0.53
Coefficient of variation	0.33	0.31	0.07
Min	17.82	7.10	39.76
Max	50.68	27.7	56.80
2012			
Income and consumption shares (129)			
Mean	0.31	0.18	0.52
Coefficient of variation	0.24	0.23	0.07
Min	20.76	6.74	30.67
Max	59.86	24.51	57.27

(continued)

Table 1 (continued)

	Richest 10%	Poorest 40%	Middle 50%
Consumption shares only (92)			
Mean	0.31	0.18	0.51
Coefficient of variation	0.23	0.22	0.08
Min	20.98	6.74	30.67
Max	59.86	24.51	55.83
Income shares only (45)			
Mean	0.30	0.17	0.53
Coefficient of variation	0.24	0.26	0.06
Min	20.76	9.15	45.20
Max	45.67	24.44	58.22

Note: Number of countries in parentheses.
Source: Authors' estimates based on World Bank (2015).

Table 2
Temporal stability of income shares 1990–2012

	Richest 10%	Poorest 40%	Middle 50%
Equal weight to every observation			
Income and consumption shares			
Grand mean**	0.32	0.17	0.52
Coefficient of variation	0.25	0.29	0.07
Min	17.14	5.70	30.67
Max	59.86	29.90	59.49
Consumption shares only			
Mean	0.30	0.18	0.52
Coefficient of variation	0.21	0.22	0.06
Min	17.14	5.78	30.67
Max	59.86	29.90	56.93
Income shares only			
Mean	0.33	0.16	0.51
Coefficient of variation	0.27	0.36	0.08
Min	17.82	5.70	38.08
Max	52.56	27.70	59.49
Equal weight to every country			
Income and consumption shares			
Mean of means***	0.31	0.17	0.51
Coefficient of variation	0.23	0.25	0.07

(continued)

Table 2 (continued)			
	Richest 10%	**Poorest 40%**	**Middle 50%**
Min	20.58	7.85	36.93
Max	55.19	25.01	55.67
Consumption shares only			
Mean of means	0.32	0.17	0.51
Coefficient of variation	0.21	0.22	0.07
Min	21.92	7.47	36.93
Max	55.19	25.01	55.21
Income shares only			
Mean of means	0.30	0.18	0.52
Coefficient of variation	0.29	0.31	0.07
Min	19.45	6.26	42.83
Max	47.68	25.49	56.36

Note: Combined samples including all cases. Grand mean = the sum of all observed values divided by total number of observations; mean of means = the mean of all countries' means (number of observations per country varies). Source: Authors' estimates based on World Bank (2015).

Table 3

Correlation of Palma with consumption share of D10 over various decile configurations (2012 or nearest data available)

D1	**D1-2**	**D1-3**	**D1-4**	**D1-5**	**D1-6**	**D1-7**	**D1-8**	**D1-9**
.971**	.993**	.999**	1.000**	.999**	.996**	.991**	.982**	.965**
D2	**D2-3**	**D2-4**	**D2-5**	**D2-6**	**D2-7**	**D2-8**	**D2-9**	**D3**
.998**	1.000**	.999**	.998**	.994**	.989**	.979**	.961**	.999**
D3-4	**D3-5**	**D3-6**	**D3-7**	**D3-8**	**D3-9**	**D4**	**D4-5**	**D4-6**
.998**	.995**	.991**	.985**	.974**	.955**	.996**	.992**	.988**
D4-7	**D4-8**	**D4-9**	**D5**	**D5-6**	**D5-7**	**D5-8**	**D5-9**	**D6**
.980**	.969**	.947**	.989**	.983**	.974**	.961**	.938**	.977**
D6-7	**D6-8**	**D6-9**	**D7**	**D7-8**	**D7-9**	**D8**	**D8-9**	**D9**
.967**	.952**	.927**	.957**	.940**	.912**	.924**	.891**	.861**

Note: N=92 in every cell and ** means significant correlation at p<0.01 level (2-tailed)

Source: Authors' estimates based on World Bank (2015)

Annex I: Conceptual Issues in Inequality Measurement

Measuring income and consumption inequality has a long history. There are five (or six) axioms for inequality measurement which are commonly cited (see Cowell, 2000, pp. 61–74), most notably the Pigou-Dalton (Dalton, 1920; Pigou, 1912) axiom that a transfer in income from a poorer to a richer person should register as a rise in inequality (or at least not a fall) and vice versa.

Of the available inequality measures, the Gini meets this principle and (almost) all others though fails what is known as the decomposability axiom that requires that overall inequality be related to constituent components of the distribution, such as population sub-groups, consistently. This means that if inequality rises among a population sub-group then inequality should register as a rise overall. For example, the global Gini does not unambiguously differentiate the separate contributions of within- and between-country inequality (it includes a significant 'overlap' or 'interaction' term between the within- and between-country contributions). The Theil index is fully decomposable, but as a measure of entropy it is rather less intuitive. Importantly, however, it is generally more sensitive to changes at the extreme ends of the Lorenz curve, whereas the Gini is more sensitive to changes in the middle of the distribution (see for full discussion Cowell, 2000, 2007; and Schröder, 2015). In terms of the common technical axioms, the Theil performs perfectly, and is often used as an alternative to the Gini but is not intuitive to a broader audience (Sen, 1973, p. 36).

The Gini is the most widely used measure of inequality, certainly in policy-related discussions. Despite its popularity, there are a range of more technical critiques of the Gini, and a substantial literature exists dedicated to finding technically superior measures of the frequency of distributions (see, for example, discussion in Duro, 2008; Frosini, 2012; Greselin, Pasquazzi and Zitikis 2013). In fact, as Atkinson (1973) and Sen (1973) both emphasise, despite the implicit suggestion that the axioms give some sense of 'objectivity', all indicators of inequality embody arbitrary value judgements. Atkinson (1973, p. 46, p. 66 and pp. 67–68), puts it thus:

> The conventional approach in nearly all empirical work [to compare distributions] is to adopt some summary statistic of inequality such as… the Gini coefficient – with no very explicit reason being given for preferring one measure rather than another… Summary measures such as the Gini coefficient are often presented as purely 'scientific', but in fact they explicitly embody values about a desirable distribution of income.

Atkinson (1973) demonstrates just why this matters, and how it ensures that the Gini is far from a 'neutral' measure of inequality. He first highlights that, in comparing two countries where the Lorenz curves do not intersect, we can say – and the Gini will suffice to do so – that the country with the curve closer to the line of complete equality is more equal than the other. When Lorenz curves cross, however, things become less clear. Atkinson presents the case of the UK and West Germany, for which the Lorenz curves then crossed at around 50 per cent of the population. The income share of the lowest-income 50 per cent is higher (closer to the 45-degree line) in West Germany, while that of the highest-income 50 per cent is closer to the line in the UK – but the Gini coefficient shows the UK to be less unequal. Having established the inescapability of normative judgements, Atkinson (1973) goes on to derive a mechanism to make explicit the actual preferences about inequality that are inherent in any given judgements on the comparison of two distributions. At a level of theory there is little to add to this. However, the complexity of Atkinson's 'equally distributed equivalent measure' approach may explain its broad absence from policy discussions

in the subsequent four decades – and this raises a further issue for measurement related to policy.

The extent to which any measure can lead or improve accountability relates to its clarity to both a policymaker and a public audience. One could ask whether the Gini is intuitively clear (unless at values of 0 and 1) or opaque to non-technical audiences. It may be better for policymakers to have a measure of inequality that is intuitive and explicit to non-technical audiences; perhaps even at the risk of violating some technical axioms.

Furthermore, one could ask: why measure inequality at all? Or: what is the purpose, in a given instance, of measuring inequality? One could argue that measuring inequality is motivated by a concern about income concentration, rather than about inequality *per se* (for example, because extremes of inequality can have damaging effects in terms of extreme poverty or conflict). In short, that inequality *per se* is not the issue of immediate policy concern, but rather excessive concentrations of income or consumption leading to societally damaging outcomes.

If the intention is to use such indicators in policy then one might equally well add a set of policy-based axioms for inequality measurement to the list of axioms for inequality measurement. These set of axioms should be seen as indicative only; a demonstration of the need for debate on axioms not to be solely a technical one but policy-related too. Such policy-based axioms might, for example include the following: (i) an 'Atkinson axiom': that the value judgements of using this indicator are sufficiently explicit; (ii) a 'policy-signal axiom': that it is clear what signal is being given to policymakers on the direction of change of inequality (increasing or decreasing); (iii) 'a clarity axiom': that it is clear to a public (that is, non-technical) audience what has changed; (iv) a 'policy-response axiom': that the policy response is sufficiently clear to policymakers (meaning how policies do or do not influence the indicator); and (v) 'a horizontal or groups axiom': that it is possible to capture horizontal (for example, gender and ethno-linguistic group) as well as vertical inequality in the indicator.

Indeed, one could argue that what is needed is a measure of inequality that has sufficient technical strength, but captures and presents the information in a policy-related or more accessible and intuitive way. Consistency with measures of horizontal inequality would add to the attraction of a given measure, since its presentation would not require additional explanation or complication.

In sum, at an analytical and policy level, it is important to make underlying judgements about inequality more explicit in measurement. For policymakers and for public discussion of inequality, it is also necessary that the chosen measure/s of inequality be easily understood and intuitively clear, as well as having clear implications for policy.

From this discussion of technical and policy axioms, one may conclude that no single measure is likely to meet every concern. As such, policy frameworks should perhaps avoid seeking single measures of inequality on which to rely entirely.

Annex II: Further Tests

Just like any measure of inequality, all measures of statistical dispersion have their strengths and weaknesses. One might object to the use of the coefficient of variation (CV) to assess the homogeneity of decile income shares across countries. The argument is that the CV is linked to the mean which itself is dependent on the design of decile groupings. In response to this one can consider alternative measures of variability to check if the CV yields a result not supported by other metrics. The Median Absolute Deviation (MAD) and the Inter-Quartile Rage (IQR), unlike the CV, are not linked to the population mean. Both metrics are independent of sample size, and are considered among the most robust, efficient and reliable measures of statistical dispersion. The MAD is defined as the median of the absolute difference between the values of a series and the series' median. The IQR is defined as the range between the upper (Q3) and lower (Q1) quartile cut-offs. Common, non-robust measures of scale can be misleading, as they might give disproportional weight to the tails of a distribution (e.g. the standard deviation which squares distances from the mean and is thus highly sensitive to outliers). Table A1 shows that both MAD and IQR support the finding that the income and consumption capture of the 'middle' 50% is considerably more homogenous across countries than that of the other decile groups. Palma's proposition is thus reproducible across different metrics of dispersion and does not depend on the use of the CV. The MAD and IQR, however, deviate from the CV when comparing the variability of the top and bottom shares. While the CV suggests that the top and bottom are, more or less, equally variable across countries, MAD and IQR suggest – plausibly, perhaps – that the top share is more variable than the bottom share.

As Table A1 shows, even the standard deviation (SD) gives us a fairly similar overall picture. The exception is 2012 consumption data which suggests greater variability of the middle than the bottom

Table A1
Measures of dispersion for cross-country decile group shares

	Consumption			Income		
	Richest 10%	Poorest 40%	Middle 50%	Richest 10%	Poorest 40%	Middle 50%
	2012 (or closest available)					
Mean	0.31	0.18	0.51	0.30	0.17	0.53
CV	0.23	0.22	0.08	0.24	0.26	0.06
SD	7.21	3.84	3.90	7.19	4.49	3.19
MAD	3.54	2.99	1.53	4.34	4.23	0.89
IQR	6.96	5.64	3.18	11.02	7.72	3.57
	1990 (or closest available)					
Mean	0.34	0.16	0.50	0.28	0.19	0.53
CV	0.23	0.29	0.08	0.33	0.31	0.07
SD	7.87	4.67	3.80	9.13	5.85	3.76
MAD	6.14	3.75	2.67	5.20	4.38	1.01
IQR	13.03	6.88	4.72	11.72	9.29	2.59

Source: Authors' estimates based on World Bank (2015).

share. What is the reason for this? If we consider the frequency distribution of D5-9 consumption data in 2012, we find that it has a heavy left tail (skewness: -2.6; kurtosis: 9.4). Taking a closer look, the SD of this series turns out to be dominated mainly by one outlier: the Seychelles, a small-island economy with a questionable data point in the year 2006 (which, moreover, is not a genuine consumption data record but a regression-based imputation provided by the PovCal database). Returning to Figure 2 in this paper it is easily visible that the Seychelles' line stands out from the rest of the data. Taking this into consideration, the SD is a misleading summary statistic for the dispersion of this series and removing the Seychelles alone would result in the SD of D5-9 to be the lowest in this year and data source as well. Still, the distribution remains skewed and we do not consider the SD a reliable measure even after removing the Seychelles. To sum up, robust measures of scale, unlinked to the population mean, confirm what was also suggested by the CV: the share of D5-9 is considerably less variable across countries than that of D1-4 and, particularly, that of D10. What about the constancy of the D5-9 consumption or income shares across points in time, though? What would an appropriate test of this constancy look like? We argue that the question remains one of descriptive statistics and that a comparison of the variation of shares is sufficient to determine if the middle is still relatively more variable than the other decile shares. Cobham & Sumner (2013b) do this. In addition to that, we can examine the frequency and distribution of changes of shares and consider measures of the central tendency and variability of change, which also show that there is little change overall and, particularly, in the share of the middle over time. Hazledine (2014) runs a regression analysis. A regression of changes of shares shows us, based on the assumption of a linear model, what change of the share of D5-9 to expect given a change in the share of either D10 or D1-4. The results suggest that any change in the income share of D10 will be explained about equally by a change in the share of D1-4 and a change in the share of D5-9. The relationship between changes of D1-4 and D5-9 is quite a bit weaker with an r^2 of around 0.2 using consumption data from 1990 (or nearest) and 2012 (or nearest). More importantly, though, regression analysis tells us nothing about the likelihood of change occurring. Yes, if there is a great deal of change in the share of D10 then there will be a corresponding (moderate) change in the share of D5-9. As descriptive statistics show: such substantial change is empirically rare. In sum, the observation that the middle shares have shown a high degree of cross-country homogeneity remains defensible even after adopting an inter-temporal perspective and using various metrics of dispersion.