



Working Paper Series

Is it all about the tails? The Palma measure of income inequality

Alex Cobham
Andy Sumner

ECINEQ WP 2013 - 308

Is it all about the tails? The Palma measure of income inequality*

Alex Cobham[†]

Centre for Global Development (CGD)

Andrew Sumner

King's College

Abstract

The “Palma” is the ratio of national income shares of the top 10 percent of households to the bottom 40 percent, reflecting Gabriel Palma’s observation of the stability of the “middle” 50 percent share of income across countries so that distribution is largely a question of the tails. In this paper we explore the Palma and corroborate the findings that the middle does indeed hold over time and through various stages of tax and transfers. Further, we find that the Gini is almost completely “explained” by only two points of the distribution: the same income shares which determine the Palma. It thus appears that both the Gini and the Palma, in practice, summarize the same information about the income distribution: but only in the case of the Palma is this explicit. This, we argue, makes the Palma a more useful (and intuitive) measure of inequality for policymakers and citizens to track.

Keywords: Inequality, Gini coefficient, Palma.

JEL Classification: D63, I3.

*We are grateful for outstanding research assistance from Alice Lépiessier, and for the assistance of Chenyue Zhao on an earlier version and the support of Save the Children, where Alex Cobham was employed during some of the research. We also acknowledge helpful comments on earlier versions of this paper from an anonymous reviewer, Roberto Angulo, Tony Atkinson, Valentina Calderon, Michael Clemens, Andrea Cornia, Stefan Dercon, Lars Engberg-Pedersen, Martin Evans, Nick Lea, Nora Lustig, Rolf Luyendijk, Richard Manning, Branko Milanovic, Nuria Molina, Chris de Neubourg, Gabriel Palma, Frances Stewart and Jan Vandemoortele; and participants in a National Development Planning Commission workshop in Accra, an Overseas Development Institute (London) seminar, a UNDP webinar, the Royal Statistics Society annual conference and the post-2015 thematic consultation on inequality. Views expressed are ours alone.

[†]**Contact Details:** Alex Cobham, Centre for Global Development (CGD), Washington, DC. E-mail: acobham@cgdev.org. Andrew Sumner, King’s College London Andrew.sumner@kcl.ac.uk.

Acronyms

D	decile
GNI	gross national income
LIC	low-income country
MIC	middle-income country
Q	quartile
OLS	ordinary least squares
SEDLAC	Socio-Economic Database for Latin America and the Caribbean
UNDP	United Nations Development Programme
WIDER	World Institute for Development Economics Research
WIID	World Income Inequality Database

1. Introduction

There are normative or instrumental reasons why inequality may be said to matter (e.g. fairness and meritocracy). However, much global literature has taken an instrumentalist approach as to why high or rising inequality can hinder development. For example, Birdsall (2007) argues that income inequality in developing countries matters for at least three instrumental reasons: where markets are underdeveloped, inequality inhibits growth through economic mechanisms; where institutions of government are weak, inequality exacerbates the problem of creating and maintaining accountable government, increasing the probability of economic and social policies that inhibit growth and poverty reduction; and where social institutions are fragile, inequality further discourages the civic and social life that underpins the effective collective decision-making that is necessary to the functioning of healthy societies.

In fact, there is empirical research that high or rising national income inequality can have a negative effect on the rate of economic growth or the length of growth spells (see, for discussion, Berg & Ostry, 2011; Cornia et al., 2004; Easterly, 2002) and high or rising national income inequality is likely to be a drag on poverty reduction (see, for discussion, Fosu, 2011; Misselhorn and Klasen, 2006; Ravallion, 2005) so while it may be the case that growth (still) is good for the poor in a general sense or at least the poorest 40% (see Dollar et al., 2013), growth is likely to be better for the poor in countries with lower initial income inequality or where income inequality is declining than in countries where the opposite is true.

In this paper we explore one particular measure of income inequality (concentration), the Palma. The Palma is a particular specification within a family of inequality measures known as ‘inter-decile ratios’, of which the most commonly used is possibly the ‘bottom 20%/top 20%’ or its inverse. The Palma is the ratio of national income shares of the top 10% of households to the bottom 40%, reflecting Gabriel Palma’s (2006, 2011) observation of the stability of the ‘middle’ 50% share of income across countries so that distribution is largely a question of the tails.

This paper sets out to assess the Palma and national income inequality trends in developing countries.² In section 2 we consider a range of axioms for inequality

² We focus our paper on developing countries (meaning low and middle income in the World Bank classification). An important caveat is that it is not as clear that the capture of the middle classes will always hold for higher-income countries. For example, the middle three quintiles (not Palma’s middle five deciles) in the USA have seen their share of national income fall from 53.2% to 45.7% between 1968 and 2011 (Levine, 2012). In the UK, the share of the middle five deciles declined only very gradually: from 56.6% in 1977 to an average of 55.6% in the 1980s, 54.7% in the 1990s, and 53.7% in the 2000s. Recent UK political discussion of a ‘middle-class squeeze’ is more likely to reflect shorter-term dynamics, with the financial crisis followed by a relatively

measures, and consider the Gini, the Theil and the Palma. We note the finding of Atkinson (1973) and Sen (1973) that any measure of inequality reflects a normative prioritisation, and that the Palma offers the potential advantage of being explicit about the prioritisation that is made.

In Section 3 we present evidence on the robustness of Palma's 'middle class capture' over time, and across stages of the income distribution, adding to the case for the use of the Palma over other inter-decile ratios. We also explore the characteristics of the Palma further, including its relationship with the 'middle' deciles' share of income, and consider patterns over time.

Section 4 addresses the relationship between the Palma and the Gini, and in particular the very high correlation. In effect, the top 10% and bottom 40% income shares almost perfectly define both the Gini and the Palma: but only in the latter case is this relationship explicit, and only in the latter case does the resulting measure lend itself clearly to policy prioritisation. We conclude that there is a strong case for the Palma to, at a minimum, sit alongside the Gini in tracking inequality.

sharp drop from 54.4% in 2008–09 to 52.9% in 2010–11 (our calculations from the UK Office for National Statistics data). Further research should consider whether there is evidence for longer-term 'middle-class squeezes', in some high-income countries in particular. In general, however, globalisation appears to be creating a distributional scenario in which what really matters is the income share of the rich (because the rest 'follows' as Palma argues).

2. Axioms for Measurement and Axioms for Policy

Measuring income inequality has a long history (for a short review of the range of inequality measures, see Charles-Coll, 2011). The Gini coefficient, for example, was developed by Corrado Gini a hundred years ago, although there are a number of other measures, such as the Theil index and inter-decile ranges (of which the Palma is a variant).

There are a set of well-known axioms for the measurement of inequality. However, there is no agreement on the exact set of axioms (what kind of global body could, after all, legitimately generate such a list?). There are five axioms for inequality measurement which are commonly cited (see Cowell, 2000, pp. 61–74). Litchfield (1999) expresses these as follows:³

- 1) the Pigou-Dalton transfer principle rules out counter-intuitive responses to transfers, e.g. the measure should not rise after a transfer from a rich person to a poor one;
- 2) income scale independence, so the measure should not respond to proportional changes in each person's income;
- 3) Dalton's principle of population, so the measure should not respond to a merging of identical populations;
- 4) anonymity or symmetry, so the measure is independent of any non-income characteristic of individuals; and
- 5) decomposability, so that (broadly) overall inequality is related consistently to inequality among sub-groups.

Of the available inequality measures, the Gini is the more widely used, arguably because of its close and relatively intuitive association, for a technical audience, with the Lorenz curve. The Gini though is not without problems. 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 discussion in Duro, 2008; Frosini, 2012; and Greselin et al., 2013).

³ Similarly, Charles-Coll (2011, p. 46) notes: The transfer principle, also known as the Pigou-Dalton principle (Dalton, 1920 and Pigou, 1912), where a transfer from a poor individual to a richer one should translate into an increase in the measure of inequality, no matter the size of the transfer or the relative position of the poor regarding the rich... The scale independence, which states that if the general income level increases by a fixed amount, then the overall value of the inequality measure should not change at all...The anonymity principle, by which the identity of the income recipients does not matter for the value determination of the inequality measure...The population independence, which means that the inequality measure should not be influenced by the size of the population.

One such issue is that the Gini is not decomposable. 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).

In terms of the common technical axioms listed, the Theil performs perfectly, and is often used as an alternative to the Gini. The problem with the Theil, as expressed by Amartya Sen, is more fundamental: it is: ‘an arbitrary formula, and the average of the logarithms of the reciprocals of income shares weighted by income shares is not a measure that is exactly overflowing with intuitive sense.’ (Sen, 1973: 36).

In fact, as Atkinson (1973) and Sen (1973) both emphasise, despite the axioms above suggesting some sense of ‘objectivity’, all indicators of inequality embody arbitrary value judgments. Atkinson (1973, p.46 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... [W]ithout introducing [judgements about the level of inequality considered ‘fair’] it is impossible to measure the degree of inequality. That no such decision has to be made with the conventional measures simply obscures the fact that they embody quite arbitrary values about the 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 United Kingdom and West Germany, for which the Lorenz curves then crossed at around 50% of the population. The income share of the lowest-income 50% is higher (closer to the 45-degree line) in West Germany, while that of the highest-income 50% is closer to the line in the UK – but the Gini coefficient shows the UK to be less unequal. Atkinson concludes:

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 (p.66).

Having established the inescapability of normative judgments, Atkinson (1973) goes on to derive an elegant mechanism to make explicit the actual preferences about inequality that are inherent in any given judgement 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 unclear (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.

One could ask: why measure inequality at all? Or: what is the purpose, in a given instance, of measuring inequality? Is it motivated by a concern about income concentrations, rather than about inequality *per se* (for example, because extremes of inequality can have damaging effects in terms of extreme poverty or conflict)? Is it because we care if standards of living differ by, for example, gender or ethno-linguistic group, or by region, age or disability? (See e.g. Stewart, 2002; Cobham & Hogg, 2010; and Kabeer, 2010.) In short, one could argue that inequality *per se* is not the issue of immediate policy concern, but rather excessive concentrations of income 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 five axioms for inequality measurement. Such policy-based axioms might include the following five:

1. An Atkinson axiom: That the value judgements of using this indicator sufficiently explicit.
2. A policy-signal axiom: That it is clear what signal is being given to policymakers on the direction of change of inequality (improving or worsening).
3. A clarity axiom: that it is clear to a public (ie non-technical) audience what has changed and what it means.
4. A policy-response axiom: that the policy response is sufficiently clear to policymakers (meaning how policies do or do not influence the indicator).
5. A horizontal or groups axiom: that it is possible to capture horizontal (e.g. gender and ethno-linguistic group) as well as vertical inequality in the indicator.

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. 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 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 addition, a given measure would be more attractive if it could say something to policymakers and public audiences about ‘wrong directions’ for inequality – subject to the implied value judgement being explicit. For example, with inter-decile ranges such as the Palma, a growing divergence between each decile’s capture of GNI and population share might be taken as a statement on ‘wrong directions’ that need policy redress.

In sum, at an analytical and policy level, it is important to make underlying judgements about inequality explicit. 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 brief discussion of potential technical and policy axioms, one can 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. One can identify serious concerns about the dominant use of the Gini as a common single measure. The Gini is oversensitive to the middle of the distribution and consequentially less sensitive to changes at the extreme. Does it matter that this is not explicit? What does one care about – the distribution in the middle or at the extremes? 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 undersensitive to the part of the distribution where change is more likely, and which we might be more concerned about – and on top of this, that the measure in question does not make these normative judgements explicit.

3. The Palma ratio

Palma (2006; 2011) observed a startling stylised fact across countries which is the capture of half of GNI by the ‘middle classes’ – defined as the five ‘middle’ deciles (deciles 5 to 9) between the extremely poor (deciles 1 to 4) and the rich (decile 10, the richest decile).⁴ On that basis, one could argue that half of the world’s population (the middle and upper-middle classes) have 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’ and the ‘extremely poor’.

Table 1 shows, for illustration, the Palma and Gini measures for the five most unequal and five most equal countries based on the World Bank’s PovCal dataset and surveys for 2010.⁵ Broadly, a Palma ratio of 1 is consistent with a Gini coefficient of around 28; a Palma of 2 a Gini of around 41; and a Palma of 3 a Gini of around 48.

The comparison to the more oft used Gini coefficient illustrates that if Palma’s findings are robust, the Palma ratio of income shares of the top 10% to the bottom 40% will capture substantial information about comparative income inequality in a single number that – arguably – is more understandable to a wider audience than the Gini.

Using a World Bank *World Development Indicators* dataset that includes observations for 135 countries with information on Gini coefficients and income shares, Palma (2011) established the claim that there are two opposite forces at work on distributions: one ‘centrifugal’, leading to an increased diversity in the shares of the top 10% and bottom 40%, the other ‘centripetal’, leading to a growing uniformity in the income share appropriated by the ‘middle’ 50% (deciles 5 to 9).

⁴ Palma here uses ‘middle class’ to mean the middle income/consumption groups. One cannot, of course, conflate social identity and expenditure data in more than the most general sense and indeed in some countries the ‘poor’ will be in the middle deciles. However, there is some basis in that the \$2 poverty rate in the middle-income countries is around 40% of population (weighted mean, all MICs) so in all but the remaining 36 LICs, the bottom four deciles is not an unreasonable proxy for the \$2 poor (Sumner, 2012). Palma (2011: 102) argues 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.’ Palma also notes a clear difference between the GNI capture of D5–D6 versus D7–D9 and a very large difference between D9 versus D10 capture of GNI.

⁵ Annex 1 contains Palma and Gini values for the sample of 79 countries used here.

Table 1: Palma and Gini for selected high and low inequality countries, 2010

Country	Palma	Gini	Year	Country	Gini	Palma	Year
The 5 most unequal on the Palma				The 5 most unequal on the Gini			
Jamaica	14.67	0.66	2002	Jamaica	0.66	14.67	2002
South Africa	7.05	0.63	2008	Namibia	0.64	7.05	2008
Namibia	6.69	0.64	2003	South Africa	0.63	6.69	2003
Honduras	5.21	0.57	2009	Zambia	0.57	5.21	2009
Bolivia	4.85	0.56	2008	Honduras	0.57	4.85	2008
The 5 least unequal on the Palma				The 5 least unequal on the Gini			
...				...			
Belarus	1.12	0.30	1998	Pakistan	0.30	1.12	1998
Kazakhstan	1.07	0.29	2009	Kazakhstan	0.29	1.07	2009
Ukraine	1.05	0.29	1999	Ukraine	0.29	1.05	1999
Bulgaria	1.00	0.28	2007	Bulgaria	0.28	1.00	2007
Romania	0.95	0.27	2011	Romania	0.27	0.95	2011

Source: PovCal.

If this observation is robust, then the ratio between the shares of the top 10% and the bottom 40% should capture the central feature of comparative income inequality. This section therefore explores the robustness of Palma's central stylised fact, specifically the stability of the middle 50%'s share of income.

First, we use decile data on income distribution from the World Bank's PovCal dataset (downloaded April 2013). We take data for the nearest dates to 1990 and 2010 for each of the 79 for which data is available for both points, subject to the following conditions: data before 1986 are excluded; we require a minimum span between the two points of ten years; and for each country, the data for both points must relate to the same survey basis (i.e. either consumption or income, to avoid comparing one with the other).

Annex I provides a table of Gini and Palma values for the selected countries, along with survey type and survey years. Around 60% of PovCal distribution data, and 70% in our sample, are consumption surveys. The remainder are income surveys (largely Latin America and the Caribbean). Because no means of adjustment (income vs. consumption) is readily acceptable we do not adjust surveys, but as noted only consider country changes by looking at surveys of the same type. In addition, we have looked at the results separately and report these where they differ significantly.

We confirm that Palma's finding of the stability of the middle 50% holds over time. Figures 1 and 2 show the income shares of the middle 50% (in green), the bottom 40% and the top 10%, for 1990 and 2010 respectively. The middle class share ranges, among the 158 observations (79 for each period), between 30.7 and 56.3 (Namibia and Guinea, respectively, in 1990); but eight out of ten observations (nine out of ten in the most recent surveys) are within the range 45%–55%. The top 10% share in contrast ranges between 19% and 65% (Belarus and Namibia, respectively), and the lower 40% share between 3% and 25% (Jamaica and Belarus).

The visual impression that the stability has increased over time is confirmed by the coefficients of variation, shown in Table 2. The 'middle class' share varies consistently less across countries than do the shares of the top 10% and bottom 40%; all three are more stable across countries in 2010 than in 1990, but the middle class has a coefficient of variation which is consistently less than a third of that of the top 10%, and around a quarter of that of the bottom 40%.

When we consider the coefficients of variation according to the type of survey (income versus consumption), the stability of the middle class is still confirmed. All income shares also tend to vary less in 2010 than in 1990. However, in income surveys the shares of the lowest 40% tend to vary twice as much as in consumption surveys, and the coefficient of variation of the lowest 40% share of income is five times greater than for the middle class and twice as big as for the top 10%. This is unlikely to be due to the smaller number of observations for income surveys (22 countries as opposed to 57) as the other coefficients of variation (CV) across survey type are more or less equal. Likewise, when we compare this to the CV on the lowest 40% for four countries based on the SEDLAC data, the CV on the lowest 40% based on income surveys is about four to five times greater. Further research may be warranted to understand better this phenomenon.

We next present additional evidence using income (rather than consumption) distribution data for countries in Latin American and the Caribbean, drawn from SEDLAC (compiled by CEDLAS and the World Bank), and for the UK (from the Office of National Statistics). This gives substantially better coverage over time than PovCal, allowing us to assess the stability of the Palma 'middle' within countries.

Figure 1: Income shares, 1990




Figure 2: Income shares, 2010




Table 2: Stability of income shares

	On the basis of income and consumption (79 countries)			On the basis of consumption alone (57 countries)			On the basis of income alone (22 countries)		
	Highest 10%	Lowest 40%	Middle 50%	Highest 10%	Lowest 40%	Middle 50%	Highest 10%	Lowest 40%	Middle 50%
Base year									
Mean Coefficient of variation	0.34	0.16	0.51	0.32	0.17	0.51	0.39	0.12	0.49
	0.26	0.34	0.09	0.23	0.26	0.07	0.25	0.46	0.11
Current year									
Mean Coefficient of variation	0.33	0.16	0.51	0.31	0.18	0.52	0.39	0.12	0.49
	0.22	0.27	0.07	0.19	0.19	0.06	0.18	0.33	0.08
Combined									
Mean Coefficient of variation	0.33	0.16	0.51	0.31	0.17	0.51	0.39	0.12	0.49
	0.23	0.29	0.07	0.20	0.21	0.06	0.21	0.38	0.09

Table 3 shows the UK and the 14 Latin American countries for which there are ten or more observations for national income distribution in SEDLAC, between 1981 and 2011. For the most part these are annual data, although in some cases they are more frequent (e.g. in Argentina they are six-monthly for some of the period).

We use again the coefficient of variation (the ratio of the standard deviation to the mean) as a measure of the stability of series. For each country individually, and across the pooled country averages, a clear pattern emerges: the income share of the middle 50% is consistently much more stable than the rest of the distribution. In general (though not without exception) as Figure 3 shows graphically, this reflects an underlying pattern that the stability of decile shares is higher for each of deciles 5–9 than for deciles 1–4 or 10. The stability is especially marked for deciles 8 and 9.

We also examine the stability of the income distribution as policy measures take effect. Here we combine data from Lustig et al. (2012), who analyse the effects of taxes and transfers for a number of Latin American countries, with Office of National Statistics data for the UK which also shows this. Table 4 shows in summary the evolution of the shares of national income of the bottom 40%, top 10% and middle 50% for Argentina, Brazil, Mexico, Peru and the UK, at three stages: market (or ‘original’) income, disposable income (i.e. market income after deductions of income tax and employees’ social security contributions, and the receipt of direct transfers) and final income (i.e. disposable income after deductions of indirect taxes, co-payments and user fees – for e.g. health care, and receipt of indirect subsidies and in-kind benefits such as public health and education).

It is clear that, even in very different countries the middle 50% share of national income is relatively untouched by systems of taxation and transfers – while the top 10%, and above all the bottom 40% are significantly affected, as is the Palma ratio. Though less relevant here, it is interesting to note that there is also strong support for the view that Latin American countries have, as yet, been unable to achieve significant redistribution through direct taxation and transfers – whereas in the UK this is responsible for the majority of redistribution.


Table 3: Stability of the ‘middle’ 50% income share in Latin America and the UK

	Obs	Year		Average values				Coefficient of variation (%)			
		Earliest	Latest	Bottom 40%	Middle 50%	Top 10%	Palma ratio	Bottom 40%	Middle 50%	Top 10%	Palma ratio
Argentina	18	2003	2011	12.2	52.5	35.2	2.93	10.4	3.2	8.1	19.2
Brazil	26	1981	2009	8.9	45.0	46.1	5.23	9.0	2.8	4.2	13.1
Chile	10	1987	2009	11.0	44.6	44.4	4.05	6.1	1.4	2.6	8.0
Colombia	13	1996	2010	9.7	44.2	46.0	4.77	4.9	2.1	2.4	7.2
Costa Rica	23	1989	2010	13.2	51.1	35.7	2.72	5.3	3.1	6.2	11.4
Dominican Rep.	14	1996	2010	12.3	48.5	39.2	3.20	5.3	2.4	4.3	9.2
Ecuador	13	1995	2010	10.7	46.7	42.6	4.07	12.1	4.3	7.0	18.0
El Salvador	16	1991	2010	11.8	50.5	37.7	3.26	11.5	2.2	5.0	15.6
Honduras	19	1991	2010	9.4	47.4	43.2	4.72	14.0	2.8	3.9	16.4
Mexico	12	1989	2010	11.7	47.4	40.8	3.53	7.9	2.3	4.7	12.2
Panama	16	1989	2010	9.5	48.8	41.7	4.43	8.8	1.6	2.7	10.2
Paraguay	13	1995	2010	10.1	47.1	42.8	4.30	11.0	3.0	4.5	14.6
Peru	15	1997	2010	11.2	49.0	39.8	3.61	9.8	4.4	8.0	18.1
Venezuela	16	1989	2010	14.1	52.4	33.4	2.41	9.2	2.2	7.1	15.4
UK	34	1977	2010-11	22.6	54.8	22.6	1.01	6.5	2.0	9.4	15.0

Ex. UK: Min		8.9	44.2	33.4	2.4
Max		14.1	52.5	46.1	5.2
Mean		11.1	48.2	40.6	3.8
Coeff. Var.		13.7%	5.6%	9.8%	22.2%
Incl. UK: Coeff. Var.		27.8%	6.4%	15.3%	30.0%

Source: calculations from CEDLAS and from ONS (UK), downloaded 8 March 2013.

Figure 3: Relative stability of income deciles' share of national income



Source: calculations from CEDLAS and from ONS (UK), downloaded 8 March 2013.

Table 4: Stability of the ‘middle’ 50% income share through taxes and transfers

		Bottom			Palma
		40%	Top 10%	Middle 50%	
Argentina	Market income	0.11	0.36	0.53	3.36
	Disposable income	0.13	0.34	0.53	2.51
	Final income	0.19	0.30	0.52	1.62
	Total change	73%	-17%	-3%	-52%
Brazil	Market income	0.09	0.45	0.46	5.10
	Disposable income	0.11	0.42	0.47	3.84
	Final income	0.16	0.37	0.47	2.23
	Total change	86%	-18%	1%	-56%
Mexico	Market income	0.11	0.41	0.48	3.80
	Disposable income	0.12	0.40	0.49	3.36
	Final income	0.15	0.36	0.49	2.35
	Total change	42%	-12%	1%	-38%
Peru	Market income	0.11	0.38	0.50	3.36
	Disposable income	0.12	0.37	0.51	3.17
	Final income	0.13	0.36	0.51	2.73
	Total change	16%	-6%	1%	-19%
UK	Market income	0.11	0.33	0.57	3.13
	Disposable income	0.19	0.27	0.54	1.44
	Final income	0.23	0.24	0.53	1.07
	Total change	117%	-26%	-7%	-66%

Source: calculations on data from Lustig et al. (2012) and from ONS (UK), downloaded 8 March 2013. Latin American data are for 2008 and 2009, UK data for 2010-11. ‘Final’ income data for Argentina do not include the effects of indirect subsidies and indirect taxes.

Since the Palma excludes information about the middle deciles, we examine how much information is lost in this way. As Figures 4–6 show, the Palma is quite closely correlated with the income share of the middle 50%. In fact, the linear fit shown nears 70% for 1990 and exceeds it for 2010 data when the outlier, Jamaica, is omitted. (separately and combined and log-linear fits, not shown, are marginally better). 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.

Figure 4: The Palma and middle 50% income share, 1990




Figure 5: The Palma and middle 50% income share, 2010





Figure 6: The Palma and middle 50% income share, 1990 and 2010 pooled (excludes Jamaica)



To consider if the Palma is stable over time we consider the inter-quartile movement of countries' Palma scores from 1990 to 2010. Table 5 shows the transition between quartiles over the period, with darker shading indicating deteriorating inequality, and lighter shading the reverse. Individual countries were able to move from the quartile of countries with lowest inequality (Q1) in 1990 to the higher end of the spectrum (Q3) in 2010 (China), while the Kyrgyz Republic moved from the highest quartile (Q4) to the lowest (Q1). Notwithstanding these particular cases, the graphic suggests significant 'stickiness' in inequality, despite the presence of mobility in each direction. Around a quarter of the sample saw an improvement in inequality (19 countries out of 79), and the same a deterioration (18 countries), while a little over half of the sample (42 countries) remained in the same quartile.

Within the latter group, it may be surprising to see the failure of Brazil to exit the highest inequality quartile, given the plaudits received. This is consistent, however, with Palma's (2011) view and the analysis of Espey et al. (2012), which shows that Brazil's achievement in reducing inequality has been to move from an extreme outlier position among countries, back towards the pack – but still with one of the highest inequalities of any major nation.

The cut-off points between quartiles have bunched up somewhat over the period. At the low end, a Palma of less than 1.33 was required for a country to be in the least unequal quartile in 1990, but by 2010 a Palma below 1.39 would suffice. At the high end, a Palma

exceeding 3.39 was required to be in the most unequal quartile in 1990, but by 2010 a Palma above 2.95 was sufficient. Average inequality within the sample fell, with the median Palma decreasing from 2.23 in 1990 to 1.88 in 2010, and the mean Palma from 2.84 in 1990 to 2.49 in 2010.

The changing Palma is shown in Figures 7 and 8 also. Figure 9 shows the relationship between the 1990 Palma and the subsequent absolute change in its value, with the initial value ‘explaining’ around 40% of the change; while Figure 10 shows the same for the Gini (33%). Figure 11 compares the Gini and Palma relationships, after excluding outliers. The Palma shows notably more ‘stickiness’, with the initial value ‘explaining’ some 55% compared to 35% for the Gini. Figure 8 shows the same relationship for the proportional change from 1990 to 2010, with a log-linear fit this time showing both the initial Gini or Palma ‘explaining’ around 40% of the subsequent change.

Figure 7: Palma (1990) and absolute change, 1990–2010 (linear)




Table 5: Transition between Palma quartiles

		2010			
		Q1 Palma < 1.39	Q2 Palma < 1.88	Q3 Palma < 2.95	Q4 Palma > 2.95
1990	Q1 Palma < 1.33	Bangladesh Belarus Bulgaria Burundi Croatia Egypt, Arab Rep. India Pakistan Poland Romania Tajikistan Ukraine	Albania Indonesia Lao PDR Latvia Lithuania Sri Lanka	China Macedonia, FYR	
	Q2 Palma < 2.23	Armenia Azerbaijan Ethiopia Kazakhstan Moldova Nepal	Cambodia Cameroon Jordan Niger Tanzania Tunisia Vietnam	Morocco Cote d'Ivoire Georgia Ghana Uruguay	Paraguay Bolivia
	Q3 Palma < 3.39	Mali	Burkina Faso Guinea Iran, Islamic Rep. Nigeria Russian Federation Thailand Turkey	Mauritania Philippines Malawi Madagascar Uganda Venezuela, RB Mozambique Malaysia Dominican Republic	Ecuador Rwanda Costa Rica
	Q4 Palma > 3.39	Kyrgyz Republic		Senegal Kenya Mexico Peru	El Salvador Chile Swaziland Panama Nicaragua Brazil Central African Republic Colombia Guatemala Zambia Honduras Namibia South Africa Jamaica

Figure 8: Palma (1990) and absolute change, 1990–2010 (linear), with outliers




Figure 9: Gini (1990) and absolute change, 1990–2010 (linear), with outliers




Figure 10: Palma and Gini (1990) and absolute change, 1990–2010, no outliers





Figure 11: Palma and Gini (1990) and proportional change, 1990–2010 (outliers excluded)



4. The Palma and the Gini

Below we consider the relationship between the Gini and the components of the Palma, but first we explore the relationship with the Palma further, by looking at the values for a stylised set of household decile income distributions (see Table 6). We fix the income share of the middle deciles (5–9) at 50% of national income, then calculate the shares of the bottom 40% and top 10% that are necessary to yield values of the Palma from one to ten. We then construct the synthetic Lorenz curve and calculate the associated Gini coefficients, using an adapted version of Hain (2005).

Table 6: Comparison of Palma and synthetic Gini values

Decile	Income shares (%)									
1	6.25	4.17	3.13	2.50	2.08	1.79	1.56	1.39	1.25	1.14
2	6.25	4.17	3.13	2.50	2.08	1.79	1.56	1.39	1.25	1.14
3	6.25	4.17	3.13	2.50	2.08	1.79	1.56	1.39	1.25	1.14
4	6.25	4.17	3.13	2.50	2.08	1.79	1.56	1.39	1.25	1.14
5	10	10	10	10	10	10	10	10	10	10
6	10	10	10	10	10	10	10	10	10	10
7	10	10	10	10	10	10	10	10	10	10
8	10	10	10	10	10	10	10	10	10	10
9	10	10	10	10	10	10	10	10	10	10
10	25.00	33.33	37.50	40.00	41.67	42.86	43.75	44.44	45.00	45.45
Palma	1	2	3	4	5	6	7	8	9	10
Gini	0.225	0.350	0.413	0.450	0.475	0.493	0.506	0.517	0.525	0.532


For simplicity, we hold equal the shares of deciles 5–9 and of deciles 1–4. This biases the reported Gini downwards, but to a limited extent only. For example, instead of holding the income shares of deciles 5–9 equal at 10% each, we can allow these to vary to be 6%, 8%, 10%, 12% and 14%. This effectively adds 0.04 to each reported Gini in Table 5. Similarly, we can allow the income shares of deciles 1–4 to vary in each case – so that, for example, decile 1’s share is 0.5% of national income less than that of decile 2, which in turn is 1% less than decile 3, which in turn is 0.5% less than decile 4. This adds 0.007 to each reported Gini. To give a specific example, the decile income shares for a Palma of 4 in Table 5 are 2.5% for deciles 1–4, 10% for deciles 5–9 and 40% for the top decile. We could arbitrarily vary these as discussed so the respective income shares are 1.5%, 2%, 3%, 3.5%, 6%, 8%, 10%, 12%, 14% and 40%. This would have the effect of changing the reported Gini from 0.450 to 0.497.

The insensitivity of the Gini above a certain level of inequality between the top 10% and the bottom 40% is noticeable. If the Palma increases from one to five, the Gini rises from 0.225 to 0.475; but if the Palma rises from five to ten, the Gini only increases from 0.475 to 0.532 (suggesting an exponential relationship, discussed below). As discussed, allowing for consistent variation within deciles 1–4 and deciles 5–9 would result in a

somewhat higher Gini; but importantly, no greater *variation* over this range. The Palma exhibits greater sensitivity to distributional changes (at the extremes rather than at the centre), resulting in higher specificity to inequality which may be valuable, for example, in regression analysis.

It is unsurprising that the Palma and Gini are highly correlated, given that they are measures of inequality in the same distribution. It is perhaps surprising, however, just how strong the relationship is. As Figure 12 shows, a simple exponential relationship provides a nearly perfect fit for the PovCal subsample we are using. If we exclude the outliers with Palma values more than twice the maximum shown, Jamaica (both years) and Namibia (1993), the fit rises to 0.9962.

Figure 12: The Palma and Gini relationship



This finding might appear to support continuing with the Gini as the established common inequality measure, if the Palma adds little new information.⁶ However, given that the Palma excludes the middle five deciles from consideration, and the Gini is in theory oversensitive to the middle of the distribution, the close relationship begs a question. If such a high correlation is only possible because the two measures are reflecting (exactly) the same information, does the Palma perfectly capture the middle of the distribution also; or does the Gini, in practice, fail to do so just as the Palma deliberately excludes it?

⁶ We are grateful to reviewers who made this argument, which encouraged us to explore the underlying reasons that follow.

To answer the question of how far the components of the Palma can explain the Gini, we run simple OLS regressions of the Gini on the shares of national income of the bottom 40% and the top 10% of households/individuals. Table 7 shows the results for our subsample of PovCal, first, and then for the full PovCal dataset. The results are striking: in each case, whether for our subsample or for the full dataset, the regression model is able to ‘explain’ 100% of the variation in the Gini.

Table 7: OLS results, PovCal data (calculated Gini)

Sample	Subsampl e	Subsampl e	Subsampl e	Full	Full	Full
Survey type	All	Income	Cons’n	All	Income	Cons’n
Bottom 40%	-1.846 ^{***} (-39.61)	-1.746 ^{***} (-16.83)	-1.898 ^{***} (-37.84)	-1.195 ^{***} (-336.60)	-1.196 ^{***} (-148.17)	-1.184 ^{***} (-592.51)
Top 10%	0.201 ^{***} (6.46)	0.312 ^{***} (4.41)	0.183 ^{***} (5.38)	0.581 ^{***} (267.68)	0.576 ^{***} (118.50)	0.585 ^{***} (491.03)
Constant	0.649 ^{***} (37.92)	0.590 ^{***} (15.16)	0.666 ^{***} (36.75)	0.419 ^{***} (328.71)	0.421 ^{***} (147.42)	0.415 ^{***} (588.11)
Observation s	158	44	114	826	309	517
Adjusted R- Squared	0.980	0.975	0.976	1.000	1.000	1.000

Note: t statistics in parentheses; * indicates $p < 0.05$, ** $p < 0.01$ and *** $p < 0.001$

When we break down the samples to look at income and consumption survey data separately, we find a slightly larger (negative) coefficient on the bottom 40% share of national income for income surveys, and a slightly smaller (positive) coefficient on the top 10% share. The central result is unaffected, however: that the bottom 40% and top 10% shares of income can perfectly explain the Gini.

We cannot be certain whether, or to what extent, this finding is a product of the approach taken by the World Bank to calculate Gini coefficients in PovCal. This involves synthetic Lorenz curve estimation from regression analysis of grouped data (e.g. income shares or mean incomes of population quantiles, or the share of the population in given income intervals), on the basis of the better performing of two alternatives of the Lorenz curve –

the General Quadratic (Villasenor and Arnold, 1989) and the Beta model (Kakwani, 1980).⁷

It is unsurprising that a regression relationship should be found between Gini coefficients obtained in this way and aggregated group data. It is the strength of the relationship that is surprising, given that the calculated Ginis are intended to summarise substantially more data than is used in these regressions.

Consider two possible, and not mutually exclusive, elements of the explanation. The first emphasises the actual relationships in the data. The more robust is Palma's stylised fact of the homogeneous middle, and the less the variation of distribution *within* the middle, the smaller the role of the middle in explaining variation in the resulting Gini – and the more that variation in the Gini will reflect variation between (and within) the bottom 40% and top 10%.

The other view emphasises potential weaknesses in the calculation of the Gini. The less information that is used to estimate Lorenz curves, the more likely that further aggregations of grouped data will be sufficient to predict the calculated Gini – and, perhaps, the less legitimate it is to consider such Ginis as a summary of the full distribution.

Minoiu & Reddy (2009) survey the literature on estimation from grouped data, with a particular focus on the two specifications used in PovCal, and test their performance with both true data and simulations. They find that for unimodal distributions the approaches work well (any bias is normally below 1%); but for multimodal distributions the bias can be larger and of uncertain sign. Assuming that the PovCal dataset includes observations where the true distribution is both unimodal and multimodal, it is not clear whether any such bias would affect our results here.

Shorrocks & Wan (2008) present an alternative method to those used in PovCal. They note of the latter that 'the quantile shares associated with the fitted functions can differ significantly from the reported values with which the procedure begins' (p.6), and for that reason propose a two-stage process which inserts an additional requirement: that characteristics of the synthetic sample (including group means) exactly match the reported values. Shorrocks & Wan test their approach and find a high degree of accuracy in reproducing individual data from grouped statistics. This is the approach used in the WIDER dataset, and so we repeat our analysis with the calculated Ginis presented here.

The WIDER dataset (WIID2) provides a broader sample, with more than 2,000 observations drawn from a checked and corrected version of the original collection of

⁷ Future research might consider whether the choice of specification affects the regression result.

survey data (WIID1), the Deininger & Squire database from the World Bank, estimates from the Luxembourg Income Study and Transmonee, and others. Again, we disaggregate to consider the higher quality WIDER data only (that is, we exclude what the dataset labels categories 3 and 4: those observations where both the income concept and the survey are problematic or unknown, or those observations classified as memorandum items and/or unreliable); and in each case again we consider separately data drawn from surveys of income and of consumption.

Table 8 shows the findings. Once again, the model provides a perfect fit, for both the higher quality subsample and the full dataset. The coefficients are very similar to those for PovCal data. There is again a small variation in the size of coefficients for income and consumption surveys, although this time in the opposite direction.

Finally, Table 9 shows the result of repeating the exercise with the WIDER data but using the ‘reported Gini’ rather than the calculated Gini values. The former are defined as ‘the one reported by the source or calculated by WIDER or Deininger & Squire for the old databases using POVCAL’ (WIID, 2008, p.9). These results should be treated with particular caution: elsewhere, Shorrocks & Wan (2008) query the use of Gini values reported in original sources: ‘It is possible that the published frequency table and Gini value refer to different sets of data for the same country and point of time, or that some of the numbers have been reported incorrectly’ (p.10).

A particular concern is that we are unable to distinguish between ‘reported’ Ginis of different types and sources. Nevertheless, the results broadly confirm the pattern, with similar coefficients and fit. However, the coefficients are notably less stable, while the adjusted R^2 are somewhat lower, ranging between 0.901 and 0.937 for the smallest groups (consumption surveys only), and between 0.984 and 0.993 where the sample size exceeds 1,000.

Table 8: OLS results, WIDER data (calculated Gini)

Sample	Higher quality	Higher quality	Higher quality	Full	Full	Full
	All	Income	Cons'n	All	Income	Cons'n
Bottom 40%	-1.205*** (-312.88)	-1.202*** (-289.30)	-1.246*** (-159.49)	-1.242*** (-386.31)	-1.241*** (-351.47)	-1.261*** (-216.84)
Top 10%	0.565*** (235.25)	0.566*** (218.30)	0.543*** (112.93)	0.541*** (281.62)	0.542*** (256.28)	0.533*** (157.18)
Constant	0.425*** (309.24)	0.425*** (286.03)	0.439*** (157.42)	0.439*** (391.22)	0.438*** (355.79)	0.445*** (220.82)
Observations	1293	1139	154	2130	1856	274

Adjusted R-Squared | 1.000 0.999 1.000 || 0.999 0.999 1.000
 Note: t statistics in parentheses; * indicates p<0.05, ** p<0.01 and *** p<0.001.

The weaker fit with ‘reported Gini’ values may reflect poorer quality data, if, for example, some values are not associated with the grouped data reported in the same place. Alternatively, the fit may reflect *higher* quality data: if the ‘reported Gini’ values are more likely to reflect the full, underlying distribution, and therefore contain more information than the calculated values, then a weaker fit would also be expected. We cannot say with certainty which of these explanations is more likely, and this may be a valuable avenue for further research.

Table 9: OLS results, WIDER data ('Reported' Gini)

Sample	Higher quality	Higher quality	Higher quality	Full	Full	Full
Survey type	All	Income	Cons'n	All	Income	Cons'n
Bottom 40%	-1.169*** (-60.68)	-1.152*** (-76.55)	-1.205*** (-9.38)	-1.143*** (-77.99)	-1.124*** (-87.06)	-1.260*** (-16.43)
Top 10%	0.551*** (45.88)	0.574*** (61.16)	0.313*** (3.95)	0.568*** (64.82)	0.589*** (76.11)	0.408*** (9.12)
Constant	0.421*** (61.25)	0.413*** (76.84)	0.490*** (10.67)	0.412*** (80.49)	0.403*** (89.49)	0.473*** (17.82)
Observations	1293	1139	154	2130	1856	274
Adjusted R-Squared	0.987	0.993	0.901	0.984	0.989	0.937

Note: t statistics in parentheses; * indicates p<0.05, ** p<0.01 and *** p<0.001.

Since the main values used by researchers and others are the calculated values from PovCal and WIID, the picture that emerges overall is one in which – in effect – the totality of variation in the Gini can be explained by the two components of the Palma. While the regression approaches used by both PovCal and WIDER to estimate Lorenz curves may be highly accurate to reconstruct individual data, they appear – at the least – to exacerbate an existing feature of the data, namely that they are dominated by the information contained in the bottom 40% and top 10% shares of national income. Again, further work is warranted to establish the extent to which Palma’s stylised fact of the homogeneous middle is responsible, and so the phenomenon is a genuine feature of actual distributions; and the extent to which the main methods for estimating Lorenz

curves from grouped data in effect reduce the real information contained in the resulting Gini values.


The first case suggests that – regardless of the Gini’s theoretical superiority as a measure of the full distribution – the reality of household and individual distributions of income, and the homogeneous middle in particular, are such that the Gini contains no more useful information than the Palma.

The second case would suggest that over-reliance on the Gini as the preferred single measure of inequality may have led to its being systematically generated from limited data, with the result that the measure does not on the whole contain the claimed information about the full distribution. In this scenario, the Palma may be seen as simply a more honest expression of the knowledge we do have about the distribution – with a significant weakness of the Gini exposed.

In either case, the claim that – in practice – the Gini values that are most commonly relied upon contain more information than the Palma ratio does not appear to be substantiated. Since the same two components of the income distribution effectively define both the Gini and the Palma, a choice between the two would consider the way in which the components are combined. It is not obvious why one would prefer the Gini calculation implied by the model (roughly 0.55 times the top 10% share of income, minus 1.2 times the bottom 40% share, plus 0.42), over the simplicity of the Palma ratio.

To explore the extent of the calculated Gini’s weakness as a measure of the full distribution, we also examine the extent to which it can be predicted from any other pair of points in the distribution. Figure 13 shows the adjusted R^2 for the equivalent regressions for two types of pair: first, other ‘tails’ that leave half of the distribution out (top 20% and bottom 30%, top 30% and bottom 20%, and top 40% and bottom 10%); and second, the pairs from other common concentration ratios: the top and bottom 20%; the top and bottom 10%; and the top 20% and bottom 40%. No other pair outperforms the Palma components, but in fact most are equally powerful in ‘explaining’ the Gini – which once again raises a question about the Gini’s ability to reflect the full distribution, at least when it is calculated from group income shares.

Figure 13: Fitting up the Gini: Adjusted R² for OLS regressions with various pairs of income shares, using PovCal or WIDER data



Using the stylised fact of the homogeneous middle, we can explore some rules of thumb for the Gini and the Palma.⁸ The income shares of the bottom 40 per cent, top 10 per cent and middle 50 per cent must sum to unity:

$$inc_{bottom40} + inc_{top10} + inc_{middle50} = 1$$

Palma’s stylised fact implies that the middle 50 per cent share is static, at around half of national income, so we can write:

$$inc_{bottom40} + inc_{top10} \cong \frac{1}{2}$$

We then substitute this expression into the identity for the Palma.

$$P = \frac{inc_{top10}}{inc_{bottom40}} \cong \frac{\frac{1}{2} - inc_{bottom40}}{inc_{bottom40}} = \frac{1 - 2 \cdot inc_{bottom40}}{2 \cdot inc_{bottom40}}$$

Or equivalently:

$$P = \frac{inc_{top10}}{inc_{bottom40}} \cong \frac{inc_{top10}}{\frac{1}{2} - inc_{top10}} = \frac{2 \cdot inc_{top10}}{1 - 2 \cdot inc_{top10}}$$

⁸We are grateful to Tony Atkinson for suggesting further exploration of these relationships.

This provides two rules of thumb for the Palma, for its derivation from either the bottom 40 per cent share of income, or the top 10 per cent share.

Similarly for the Gini, we can take the regression equations which we have seen define the calculated PovCal or WIDER Ginis, where a and b are the coefficients on the bottom 40 percent and top 10 per cent income shares respectively, and substitute in the stylised relationship.

$$G = a. inc_{bottom40} + b. inc_{top10} + c$$


$$= a. inc_{bottom40} + b \left(\frac{1}{2} - inc_{bottom40} \right) + c = (a - b) inc_{bottom40} + \frac{b}{2} + c$$

Or equivalently:

$$G = (b - a) inc_{top10} + \frac{a}{2} + c$$

We then use the regression results for the full PovCal sample, and the higher-quality WIDER sample, to generate the specific rules of thumb for each dataset, using either the bottom 40% income share alone, or the top 10% income share alone. To illustrate the precision of each rule of thumb estimate, we perform regressions with each rule of thumb measure as the sole independent variable, and report the adjusted R^2 in figure 14. The bottom 40% income share alone is sufficient to all but perfectly estimate the Gini in PovCal (adjusted R^2 of 0.998) and WIDER (0.996). The same share performs nearly as well in estimating the Palma in PovCal (0.99), and still strongly with the WIDER Palma (0.974). There is more of a divergence when estimating with the top 10% share only, which performs well with the Gini (PovCal 0.99 and WIDER 0.978), but very poorly with the Palma (0.102 and 0.002).

Figure 14: Rules of thumb: Adjusted R² for OLS regressions with ‘rule of thumb’ Palma and Gini based on a single income share



The Gini regression results in this section reflect the strength of its correlation with the Palma component income shares, shown in Table 10 for the full PovCal and higher-quality WIDER data respectively. It is clear that in practice the Gini – as calculated from group income shares, at least – reflects much more limited information than the full distribution.

Finally, table 11 shows the results of regressing the Gini on the component income shares separately (with no constant), which provides an alternate set of rules of thumb. With a level of accuracy indicated by the adjusted R² values, it is possible to estimate Gini values from PovCal or WIDER from just one point of the distribution, either the top 10% income share or the bottom 40% income share, simply by multiplying by the relevant coefficient. In particular, multiplying the top 10% income share by 1.268 (PovCal) or 1.278 (WIDER) estimates the Gini with 99.7% accuracy or greater.

Table 10: Correlation of group income shares with the Gini

	PovCal Calculated Gini	WIDER Calculated Gini	WIDER 'Reported Gini'
Top 10%	0.986	0.973	0.969
Top 20%	0.997	0.993	0.986
Top 30%	0.999	0.998	0.990
Top 40%	0.999	0.997	0.988
Bottom 40%	-0.991	-0.986	-0.976
Bottom 30%	-0.984	-0.972	-0.963
Bottom 20%	-0.973	-0.948	-0.938
Bottom 10%	-0.947	-0.888	-0.880

Table 11: Single income share regressions for the Gini

	Full PovCal	Full PovCal	Higher quality WIDER	Higher quality WIDER
Bottom 40%	2.249 ^{***} (47.18)		1.749 ^{***} (57.08)	
Top 10%		1.268 ^{***} (723.75)		1.278 ^{***} (662.69)
Observations	826	826	1293	1324
Adjusted R^2	0.729	0.998	0.716	0.997

t statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

5. Conclusions

The Palma is a measure of income inequality or income concentration - based on the observation of Gabriel Palma that the middle classes tend to capture around 50% of national income, so that distributional politics can be thought of, simplistically, as determining the split of the remaining half of national income between the richest 10% and the poorest 40%.

As both Atkinson (1973) and Sen (1973) made clear, it is important to recognise that no measure of inequality, including the Gini, is 'neutral': the best we can do is to be explicit about the normative decision being taken in the choice of any given measure. This is certainly true of the Palma.

We have corroborated the surprising stability of 'middle class capture' across countries, and across time, while confirming much greater variation in the Palma ratio of the top 10% and bottom 40% income shares and we have found the Palma and the Gini to have a near-perfect fit – suggesting that much of the same information is captured by the two measures. Indeed, the components of the Palma ratio alone are able to 'explain' between 99% and 100% of Gini variation. In practice, we find that no more information is contained in the Gini – a measure of the entire income distribution – than in the Palma ratio, which excludes completely the 5th to 9th deciles. Further research will be needed to evaluate the extent to which this finding is driven by Gabriel Palma's stylised fact of the 'homogeneous middle' of the distribution, and to what extent the finding results from oversimplistic calculation methods used to generate the most widely used Gini series. Even simple rules of thumb based on a single point of the distribution seem able to predict the Gini with an accuracy approaching 100%. The same holds for the Palma, if the income share of the bottom 40% is used; but in this case the finding is by construction. In the case of the Gini, the results reveal a hitherto hidden lack of depth.

We would conclude, that the robustness of Palma's thesis and the intuitive nature of the Palma ratio provide a strong case for further exploration of the Palma. We would argue that the Palma may be a better measure for policymakers and citizens to track as it is intuitively easier to understand for policymakers and citizens; that it is a more policy-relevant measure of inequality because, given the observed stability of the middle income deciles, it is clear what change is implied by a desire to change the Palma; and that it is explicit about the assumed preferences in regard to inequality.

An obvious criticism of the Palma is that it only considers half of the income distribution; for which reason we consider a measure of concentration rather than the full distribution. However, since it turns out that the Gini in practice does not capture any additional information, and moreover that it does this in an opaque and hitherto undiscovered way (as far as we are aware), we consider this an argument in favour of the Palma. Following a similar line of thought, it is worth noting that the Palma does directly expose the top

decile somewhat – which in many countries may not be appreciated – but it is the Palma’s simplicity which may be its greatest strength. 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% earn five times the income of the poorest 40% of the nation.

References

- 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 A, T. Piketty and E. Saez, 2011, Top Incomes in the long Run of History, *Journal of Economic Literature* 49:1, pp.3 – 71.
- Berg, A. & J. Ostry, 2011, Inequality and Unsustainable Growth: Two Sides of the Same Coin? IMF Staff Discussion Note SDN/11/08.
- Birdsall, N., 2007, *Income Distribution: Effects on Growth and Development*. CGD Working Paper 188. Washington DC: CGD.
- Charles-Coll, J., 2011, Understanding Income Inequality, *International Journal of Economics and Management Sciences*, 1, 3, 17-28.
- Cobham, A. & A. Hogg, 2010, *Poverty Over: We're all in this together*, London: Christian Aid.
- Cornia, G.A., Addison, T. and Kiiski, S., 2004, *Income Distribution Changes and Their Impact in the Post-Second World War Period*. In *Inequality, Growth, and Poverty in an Era of Liberalization and Globalization*. Ed. Giovanni Andrea Cornia. Oxford: Oxford University Press
- Cowell, F. A. (2000). Measurement of inequality. In A. B. Atkinson and F. Bourguignon (Eds.), *Handbook of Income Distribution*, Chapter 2, pp. 87-166. Amsterdam: North Holland.
- Cowell, F. A. (2007). *Measuring Inequality* (Third ed.). Hemel Hempstead: Oxford University Press.
- Dalton, H. (1920). The measurement of the inequality of incomes. *Economic Journal*, 30:348- 61.
- Dollar, D. Tatjana, K., Kraay, A. 2013. Growth Still is Good for the Poor. World Bank Policy Research Working Paper No. 6568. Washington DC: World Bank.
- 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.1041–1045.
- Easterly, W. (2002). *Inequality Does Cause Underdevelopment*. Centre for Global Development Working Paper Number 1. Washington, DC: Centre for Global Development.
- Espey, J., A. Holder, N. Molina & A. Cobham, 2012, *Born Equal: How reducing inequality could give our children a better future*, London: Save the Children.
- Fosu, A. (2011). *Growth, Inequality, and Poverty Reduction in Developing Countries Recent Global Evidence*. WIDER Working Paper No. 2011/01. WIDER: Helsinki.
- Frosini, B., 2012, 'Approximation and decomposition of Gini, Pietra–Ricci and Theil inequality measures', *Empirical Economics* 43(1), pp.175-197.
- Greselin, F., L. Pasquazzi & R. Zitikis, 2013, 'Contrasting the Gini and Zenga indices of economic inequality', *Journal of Applied Statistics* 40(2), pp.282-297.

- Hain, D., 2005, Gini calculator, Public Health England: Downloaded at:
<http://www.erpho.org.uk/viewResource.aspx?id=12222> (accessed 7 Sept 2013)
- HLP, 2013, A New Global Partnership: Eradicate Poverty And Transform Economies Through Sustainable Development. The Report of the High-Level Panel of Eminent Persons on the Post-2015 Development Agenda, New York: United Nations.
- Kabeer, N., 2010, Can the MDGs provide a pathway to social justice? The challenges of intersecting inequalities, New York/Brighton: UN MDG Achievement Fund and Institute of Development Studies.
- Kakwani, N., 1980, 'On a class of poverty measures', *Econometrica* 48(2), pp.437–446.
- Klasen, S., T. Otter & C. Villalobos Barría, 2012, 'The dynamics of inequality change in a highly dualistic economy: Honduras, 1991–2007', UNU-WIDER Working Paper 2012/17, Helsinki: United Nations University World Institute for Development Economics Research.
- Levine, L., 2012, 'The US income distribution and mobility: Trends and international comparisons', CRS Report R42400, Washington, DC: Congressional Research Service.
- Litchfield, J., 1999, 'Inequality: Methods and tools', World Bank:
<http://siteresources.worldbank.org/INTPGI/Resources/Inequality/litchfie.pdf>
(accessed 7 Sept 2013)
- Lustig, N., G. Gray-Molina, S. Higgins, M. Jaramillo, W. Jiménez, V. Paz, C. Pereira, C. Pessino, J. Scott & E. Yañez, 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.
- Misselhorn, M. and Klasen, S. (2006). Determinants of the Growth Semi-Elasticity of Poverty Reduction, Proceedings of the German Development Economics Conference, Berlin.
- Milanovic, B., 2005, *Worlds Apart: Measuring global and international inequality*, New Jersey: Princeton University Press.
- Minoiu, C. & S. Reddy, 2009, 'Estimating poverty and inequality from grouped data: How well do parametric methods perform?', *Journal of Income Distribution* 18(2), pp.160-178.
- Neubourg C. de (2010), *Social Security and Nation Building*, in Peter Townsend (ed.), *Building Decent Societies*, Palgrave Macmillan, Basingstoke.
- Palma, J.G., 2006, 'Globalizing inequality: 'Centrifugal' and 'centripetal' forces at work' DESA Working Paper 35, New York: UN Department of Economic and Social Affairs.
- 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', *Cambridge Working Papers in Economics* 1111, Cambridge: University of Cambridge Department of Economics (later published in *Development and Change*, 42, 1, 87-153).
- Pigou, A. C. (1912). *The economics of welfare*, London, Transaction publishers.

- Ravallion, M. 2005. Inequality is Bad for the Poor. World Bank Policy Research Working Paper No. 3677. Washington DC: World Bank.
- Sen, A. 1973, *On Economic Inequality*. New York: Norton.
- Shorrocks, A., & G. Wan, 2008, 'Ungrouping income distributions: Synthesising samples for inequality and poverty analysis', UNU-WIDER Research Paper 2008-16.
- Stewart, F., 2002, 'Horizontal inequalities: A neglected dimension of development', Queen Elizabeth House Working Paper 81.
- Sumner, A. 2012, *Where will the World's Poor Live?* CGD Working Paper 305, Washington DC: Center for Global Development.
- UNICEF & UNWomen, 2013, *Synthesis Report on the Global Thematic Consultation on Addressing Inequalities*, New York: UNICEF/UNWomen.
- Villasenor, J. & B. Arnold, 1989, 'Elliptical Lorenz curves', *Journal of Econometrics* 40, pp. 327–338.
- WIID, 2008, 'World Income Inequality Database: User guide and data sources', Helsinki: United Nations University World Institute for Development Economics Research (UNU-WIDER).

Annex I: Dataset

Country	Survey type	Base year	Gini, base year	Palma, base year	Current year	Gini, current year	Palma, current year
Albania	Consumption	1996	29.120	1.045	2008	34.510	1.436
Armenia	Consumption	1998	36.010	1.532	2010	31.300	1.220
Azerbaijan	Consumption	1995	34.960	1.433	2008	33.710	1.363
Bangladesh	Consumption	1991	27.600	0.997	2010	32.120	1.272
Belarus	Income	1988	22.760	0.757	1998	30.280	1.123
Bolivia	Income	1990	42.040	2.077	2008	56.290	4.847
Brazil	Income	1990	61.040	6.447	2009	54.690	4.302
Bulgaria	Consumption	1989	23.430	0.795	2007	28.190	0.997
Burkina Faso	Consumption	1994	50.710	3.231	2009	39.790	1.859
Burundi	Consumption	1992	33.330	1.328	2006	33.270	1.347
Cambodia	Consumption	1994	38.280	1.736	2009	36.030	1.543
Cameroon	Consumption	1996	40.680	1.944	2007	38.910	1.763
Central African Rep.	Consumption	1992	61.330	6.897	2008	56.300	4.505
Chile	Income	1990	55.250	4.235	2009	52.060	3.506
China	Consumption	1990	32.430	1.252	2009	42.060	2.080
Colombia	Income	1992	51.450	3.437	2010	55.910	4.520
Costa Rica	Income	1990	45.300	2.486	2009	50.730	3.333
Cote d'Ivoire	Consumption	1988	36.890	1.585	2008	41.500	2.026
Croatia	Consumption	1998	26.820	0.934	2008	33.650	1.356
Dominican Rep.	Income	1989	50.460	3.296	2010	47.200	2.746
Ecuador	Income	1987	50.490	3.387	2010	49.260	3.061
Egypt, Arab Rep.	Consumption	1990	32.000	1.261	2008	30.770	1.194

Country	Survey type	Base year	Gini, base year	Palma, base year	Current year	Gini, current year	Palma, current year
El Salvador	Income	1991	53.970	4.134	2009	48.330	2.951
Ethiopia	Consumption	1995	39.960	1.876	2010	33.600	1.352
Georgia	Consumption	1996	37.130	1.598	2010	42.100	2.094
Ghana	Consumption	1991	38.130	1.700	2005	42.760	2.172
Guatemala	Income	1989	59.600	5.974	2006	55.890	4.524
Guinea	Consumption	1991	46.840	2.884	2007	39.350	1.805
Honduras	Income	1990	57.360	5.013	2009	56.950	5.209
India	Consumption	1993	30.820	1.186	2009	33.900	1.392
Indonesia	Consumption	1990	29.190	1.093	2010	35.570	1.486
Iran, Islamic Rep.	Consumption	1990	43.600	2.271	2005	38.280	1.709
Jamaica	Income	1990	65.250	14.351	2002	65.700	14.669
Jordan	Consumption	1992	43.360	2.232	2010	35.430	1.489
Kazakhstan	Consumption	1996	35.320	1.442	2009	29.040	1.066
Kenya	Consumption	1992	57.460	4.735	2005	47.680	2.810
Kyrgyz Rep.	Consumption	1993	53.700	4.207	2011	33.380	1.327
Lao PDR	Consumption	1992	30.430	1.169	2008	36.740	1.599
Latvia	Consumption	1998	33.520	1.330	2009	34.810	1.417
Lithuania	Consumption	1996	32.260	1.252	2008	37.570	1.640
Macedonia, FYR	Consumption	1998	28.130	0.994	2010	43.560	2.261
Madagascar	Consumption	1993	46.120	2.572	2010	44.110	2.329
Malawi	Consumption	1997	50.310	3.163	2010	43.910	2.301
Malaysia	Income	1989	46.170	2.597	2009	46.210	2.627
Mali	Consumption	1994	50.560	3.268	2010	33.020	1.294
Mauritania	Consumption	1993	50.050	3.092	2008	40.460	1.921

Country	Survey type	Base year	Gini, base year	Palma, base year	Current year	Gini, current year	Palma, current year
Mexico	Income	1992	53.750	3.900	2010	47.730	2.812
Moldova	Consumption	1997	36.900	1.585	2010	33.030	1.299
Morocco	Consumption	1990	39.200	1.791	2007	40.880	1.958
Mozambique	Consumption	1996	44.490	2.357	2007	45.660	2.499
Namibia	Income	1993	74.330	15.081	2003	63.900	6.693
Nepal	Consumption	1995	35.230	1.482	2010	32.820	1.298
Nicaragua	Income	1993	57.280	5.130	2005	52.350	3.655
Niger	Consumption	1992	36.100	1.536	2007	34.550	1.431
Nigeria	Consumption	1992	44.950	2.463	2011	39.740	1.840
Pakistan	Consumption	1990	33.230	1.332	2007	30.020	1.156
Panama	Income	1989	58.910	6.494	2010	51.920	3.627
Paraguay	Income	1990	40.840	1.896	2010	52.420	3.730
Peru	Income	1997	53.720	4.118	2010	48.140	2.948
Philippines	Consumption	1991	43.820	2.284	2009	42.980	2.183
Poland	Consumption	1992	26.700	0.916	2011	32.730	1.282
Romania	Consumption	1998	29.440	1.078	2011	27.420	0.951
Russian Federation	Consumption	1993	48.380	2.915	2009	40.110	1.885
Rwanda	Consumption	2000	51.510	3.350	2010	50.820	3.216
Senegal	Consumption	1991	54.140	4.090	2011	40.300	1.901
South Africa	Consumption	1993	59.330	5.690	2008	63.140	7.052
Sri Lanka	Consumption	1990	32.480	1.292	2009	36.400	1.571
Swaziland	Consumption	1994	60.650	5.858	2009	51.490	3.513
Tajikistan	Consumption	1999	29.010	1.052	2009	30.830	1.155
Tanzania	Consumption	1991	33.830	1.357	2007	37.580	1.653

Country	Survey type	Base year	Gini, base year	Palma, base year	Current year	Gini, current year	Palma, current year
Thailand	Consumption	1990	45.270	2.405	2010	39.370	1.795
Tunisia	Consumption	1990	40.240	1.886	2010	36.060	1.509
Turkey	Consumption	1987	43.570	2.246	2010	40.030	1.871
Uganda	Consumption	1989	44.360	2.370	2009	44.300	2.332
Ukraine	Income	1988	23.310	0.778	1999	28.960	1.049
Uruguay	Income	1989	42.370	2.109	2010	45.320	2.470
Venezuela, RB	Income	1989	43.840	2.285	2006	44.770	2.404
Vietnam	Consumption	1992	35.680	1.508	2008	35.570	1.489
Zambia	Consumption	1993	52.610	3.897	2010	57.490	4.768

Source: World Bank PovCalNet (downloaded 30 April 2013), and our calculations.