

Empirical evidence on long-term global trends in income inequality

Joe Hasell

Our World In Data, Oxford Martin Programme
on Global Development

1. Introduction

The aim of the present paper is to say something about long-run trends in income inequality *within* countries that is representative of the movements across the *globe*. Naturally, this requires some means of aggregating over countries, considering, for instance, how the average Gini index across all countries has evolved. In Section 2 I outline the most commonly-used approach, based on national household survey data, and present some summary figures on inequality across the globe on this basis.

As we will see, conclusions about recent trends are sensitive to how we approach their measurement. Indeed, even our ability to say *whether* there has been a significant rise in inequality – looking across the globe as a whole – is dependent upon our choice of data and metric. In any case, the range of movement appears small compared to the much larger differences between countries in terms of their average incomes. Moreover, the recent between-country convergence has seemingly outpaced any increase in within-country inequality that may have occurred since the late 1980s, leading to *falling* levels of inequality between global citizens, as demonstrated in Section 3.

Section 4 presents some crucial caveats to this rather benign view of within-country inequality trends painted by the Gini coefficient, and as derived from household surveys. In particular, the approach's ability to properly capture rising top incomes has been called into question and its focus on relative, as opposed to absolute differences in incomes may also be queried. From the perspective afforded by such alternatives, the assessment of global inequality trends may be completely reversed.

In section 5 we zoom-in to consider trends in individual countries, and across country groups, revealing a high degree of heterogeneity, even among relatively similar countries. As is revealed from such a vantage point, no one story fits the experiences of all countries – a fact that gets masked by global averages. Attention to such nuances is, however, paramount. They rebut any over-simplified narrative that would have us see rising inequality as an automatic, inevitable outcome of global economic forces, and demonstrate the continued relevance of country-level policymaking for tackling inequalities.

2. Constructing a global average Gini

Whilst average incomes in a given country are easily obtained by dividing national accounts aggregates, either GDP or GNI, by the size of the population, there is no such ready-made way of seeing how that total income is actually *distributed* across those individuals. The most common way of approaching this issue is through the use of household surveys, in which a sample of respondents are asked to report their household's income, or as is more common in developing countries, on the value of the goods it consumes over a given period. To make meaningful comparisons, either in one country over time, or across different countries, requires a certain degree of homogeneity in the concepts and methods employed. We can think of these factors that may impact on comparability as a 'check-list' of questions (Atkinson and Bourguignon, 2015). Which concept of welfare is being analyzed – income, after or before taxes; consumption; wealth; non-monetary dimensions? ("inequality of what"?). What is the unit of analysis – individuals, households, families, tax units? To what extent are differences in household composition accounted for – does the data employ some method of equalization, and if so, which? Beyond these also lie a host of more technical question pertaining to the methods employed in the survey, with the specific questions being asked and the timeframe impacting the resulting data.

As might be imagined, the further our analysis of trends extends in time and space, the further from this ideal of comparability we have to travel. To say something summary about inequality across *all* countries inevitably means employing a rather mixed bag of sources, with heterogeneity along many of the dimensions above. As Lakner and Milanovic (2016) note, given the absence of a global household survey, we have no choice but to accept this lack of comparability if we wish to say something representative of inequality in the world as a whole.

The starting point when looking for survey data with global coverage is POVCAL, the World Bank's database of harmonized summary statistics derived from national household surveys across the globe. Whilst the source benefits from its wide coverage, it should be born in mind that this comes at the expense of including both income and consumption surveys. Moreover, no information is made available as to the income concept employed in the income surveys, whether net or gross, further hindering its interpretation. A number of other sources offer a more consistent harmonization across surveys, but lack the scope of POVCAL. This includes *Luxembourg Income Study* and OECD data, both of which have a focus on high-income countries, and the regionally-focused *Socio-Economic Database for Latin America and the Caribbean* (SEDLAC) and EU-SILC database, which are often resorted to in order to patch up gaps in POVCAL's coverage.¹

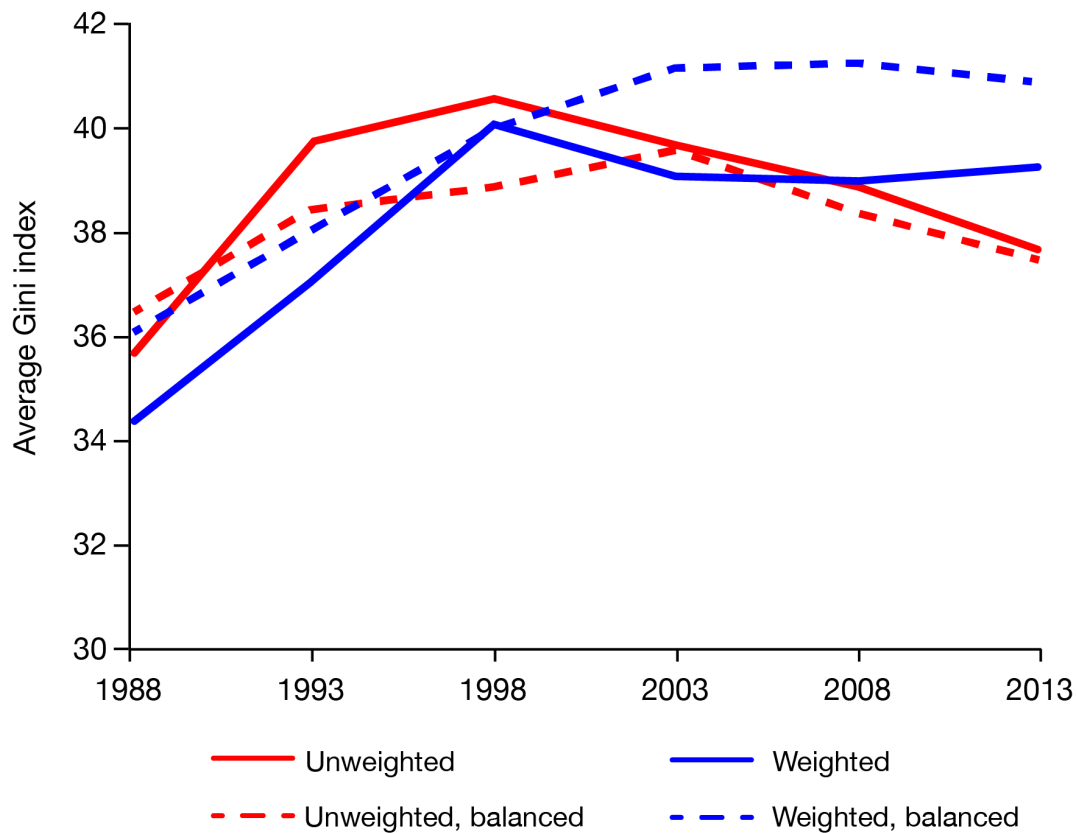
With these caveats in mind, Fig. 1 plots the global average Gini index, as presented in World Bank (2016). The figures are computed around five-year intervals from 1988 to 2013. The bulk of the Gini estimates are taken from POVCAL, and therefore juxtapose consumption- and income-based estimates.² The blue lines show the simple average across all countries, whereas the red lines show the average weighted according to the size of the population in each country. The solid lines are computed across all the countries for which a Gini estimate

¹ See, for instance, Milanovic 2014, which is employed to compliment the POVCAL data in both World Bank (2016) and Lakner and Milanovic (2016).

² The dataset also include estimates from Milanovic (2014), particularly in the case of high-income countries which are less well-represented in POVCAL.

was available at each reference period, and therefore potentially reflect compositional shifts. The dotted lines use a balanced sample, only including countries for which estimates were available for each reference year. Broadly, they corroborate the main features of the full-sample trend.

Fig. 1: Average within-country inequality, 1988–2013



Source: Reformatted from World Bank 2016. Calculation based on PovalNet and Milanovic 2014

We see in the unweighted average an increase roughly through to the millennium, followed by a fall through the 2000s that extends into this decade, finishing in 2013 only marginally higher than the level in 1988. The population-weighted average saw a somewhat more dramatic rise through the 90s, with only a levelling-out in the new century and with no fall. The difference can be explained by rising or sustained high levels of inequality seen in the 21st Century in many of the most populous countries, including China, India, and Indonesia.

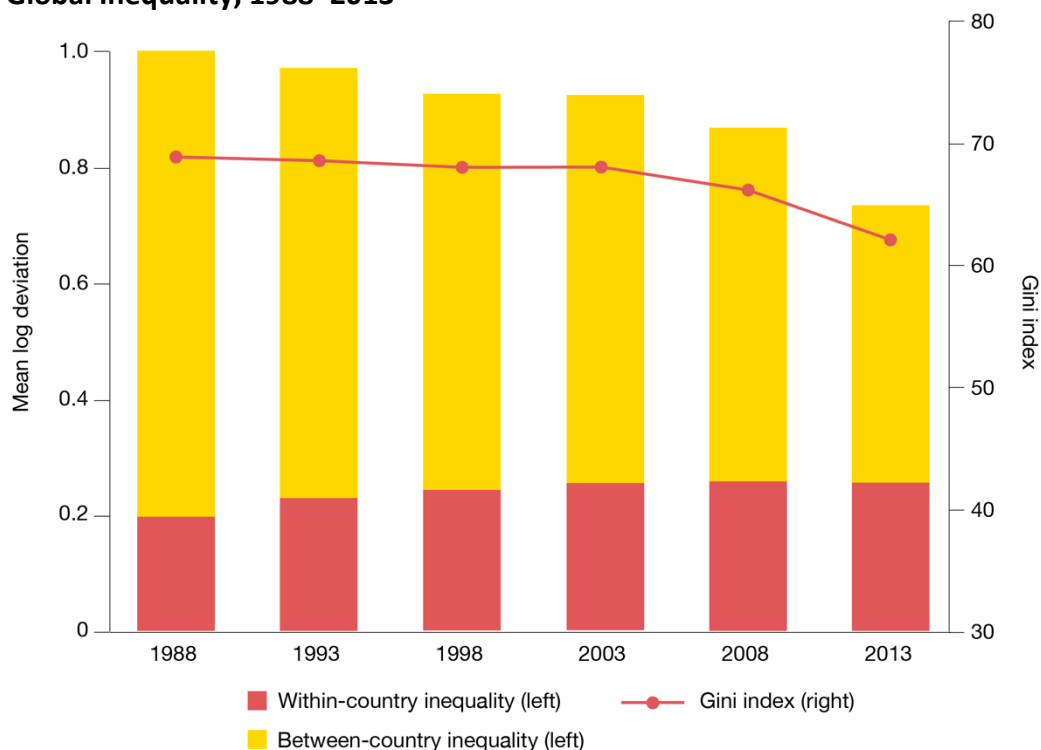
3. Within-country inequality in a global context

Given the increased prominence the issue of inequality has taken on in recent years in policy and public debate, the variation in the average Gini coefficient over time may strike us as surprisingly modest. One standard for judging the size of this shift is against movements in *global inequality* – that is to say, a summary measure of the dispersion of incomes across all individuals on the planet, irrespective of their country of residence –

using a decomposition into *between*-country and *within*-country components. Fig. 2 presents just such a composition from World Bank (2016), using the same survey data as that used in fig 1.

Since it is computed from a reconstruction of the global distribution, rather than simply averaging out country-level Ginis, the within-country component (red bars) is similar in interpretation to the weighted average Gini in Fig 1. We see that the trend is as before: increasing to the millennium, and relatively stable thereafter. As is immediately clear though, inequality within countries forms a relatively small part of total inequality across the globe, given the much larger contribution of between-country differences (yellow bars). This is increasingly less true, however, as countries' average incomes have converged due to higher growth rates among developing countries. Thus we are moving towards a world in which within-country inequality will play an increasingly important role in the global income distribution – even as it would appear to have levelled out in its own terms for the time being. The corollary of this convergence between countries having outdone the increase in inequality within countries, is that global inequality amongst all citizens taken together has fallen over this period, as measured either by the mean logarithmic deviation or Gini coefficient.³

Fig. 2: Global Inequality, 1988–2013



³ Bourguignon (2016) in fact finds a within-country component that is relatively stable until the millennium and rising thereafter. This is surprising given that both analyses make use of similar underlying data. It is not clear what in the two methodologies is driving this difference. The discrepancy can, however, be taken to emphasise the main point being made here: that whatever increase in within-country inequality that may have occurred – which would appear to be quite sensitive to methodological choice – it is relatively small compared to the between-country convergence in incomes. This is, at least, the message to be taken from household survey data – the reliability of which is interrogated in more detail in the next section.

4. Crucial caveats

The foregoing discussion has painted a fairly benign picture of inequality over the last three decades. There are, however, a number of very legitimate concerns about whether this approach is appropriately capturing the relevant changes that seem to have motivated the recent upswell in attention to the topic of inequality.

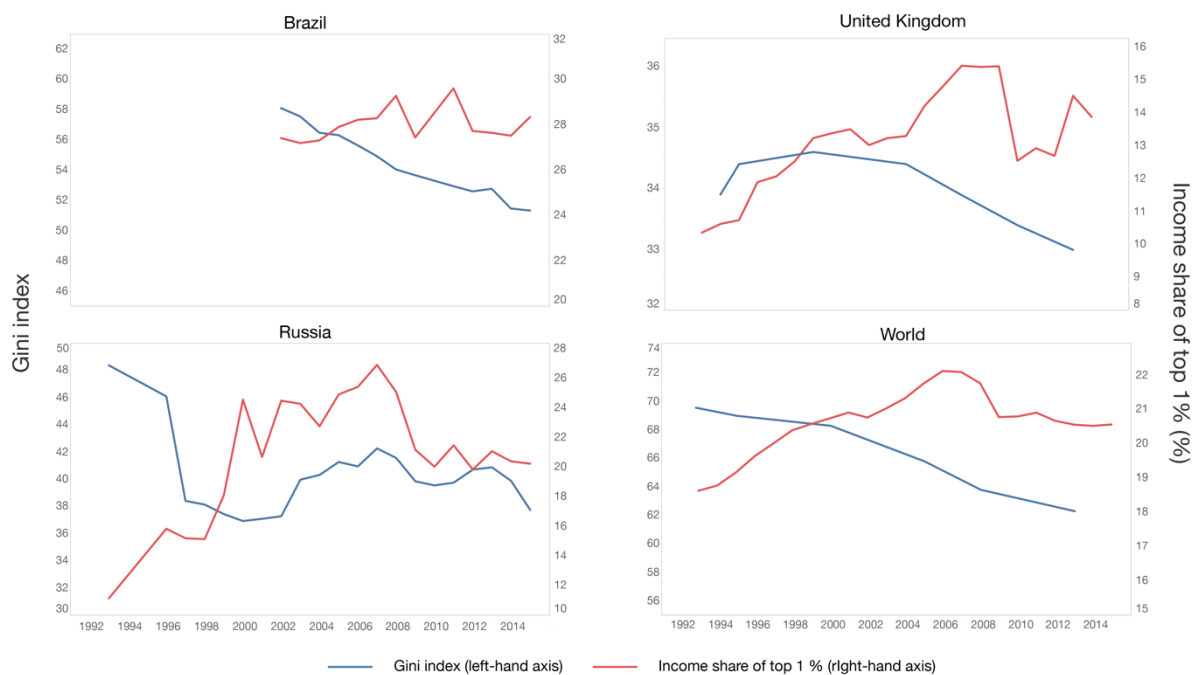
The incompatibility between the consumption and income surveys employed in global analyses has already been noted. Importantly, there is good reason to believe that consumption surveys tend to understate inequalities as compared to income, and that this bias increases with income. Richer people tend to save more of their income, and also typically spend more on durable goods which is often less-well captured in consumption surveys. Not only does this affect our ability to meaningfully compare inequality across countries and regions at a given point in time, we may also be concerned that the continued use of consumption surveys in some higher-income settings, such as is the case for urban China, may also affect trends.

Another weakness of the household survey data lies in its well-documented inability to properly capture incomes at the top of the distribution, as a consequence of richer households under-reporting their incomes or choosing not to participate in the survey at all.⁴ Again, this may be not only underestimating the level of inequality, both in individual countries and globally, it may also be understating any increases in inequality occurring at either level.

Fig. 3 makes use of the World Inequality Database series of top income shares, to compare evolution of the share of income received by the top 1% to the Gini index derived from household surveys for Russia, Brazil and UK. The same comparison for the world income distribution is also shown.

⁴ Also the small number of individuals at the very top of the distribution makes it harder to estimate top incomes with precision. This can be addressed by over-sampling the rich, but this is by no means standard practice across countries.

Fig. 3: Gini index vs income share of top 1%



Source: All top income shares data are from World Inequality Database. Gini series for Russia and Brazil are from PovcalNet, and for UK are taken from LIS Key Figures. World Gini figures are from Bourguignon 2016.

The three countries (and also the scales up the vertical axes) were intentionally selected to exemplify the possibility of rising top income shares at the same time as falling Gini of disposable income (/consumption). This is by no means always the case – indeed empirically, there is in general a close correlation between the two (Morelli, Smeeding and Thompson, 2015). Nonetheless, we see from the world estimates that a focus on the top 1% income share reverses our view of the recent global inequality trends, having risen by around 2-3 points since the early 1990s at the same time as the global Gini has fallen.

It is important to note that there are a number of differences being captured simultaneously here. The top shares data refer to pre-tax income, whereas the Gini here captures either consumption (Russia), disposable income (UK), or a mix (World).⁵ But, crucially, the approach makes use of different underlying sources – relying on administrative tax and national accounts data to calculate the ratio of top incomes to pre-tax national income. This therefore avoids the issue of under-reporting and non-response encountered in household survey data.⁶

Correcting survey data for under-reporting is possible, though there is little systematic evidence is available to document the sources of the problem as well as its extent across the

⁵ The Gini for the UK is taken from the LIS Key Figures and refer to equivalised disposable household income. The figures for Brazil are from POVCAL, and refer to income, but the specific income concept is unclear from the source. The world Gini figures are taken from Bourguignon (2016) and therefore differ slightly from the World Bank (2016) figures in Fig. 1.

⁶ Though the ability of the very richest to hide income, through offshore accounts and other structures for tax avoidance or evasion, mean that the very highest incomes may still not be fully accounted for (Alstadsæter, Johannesen, and Zucman, 2017).

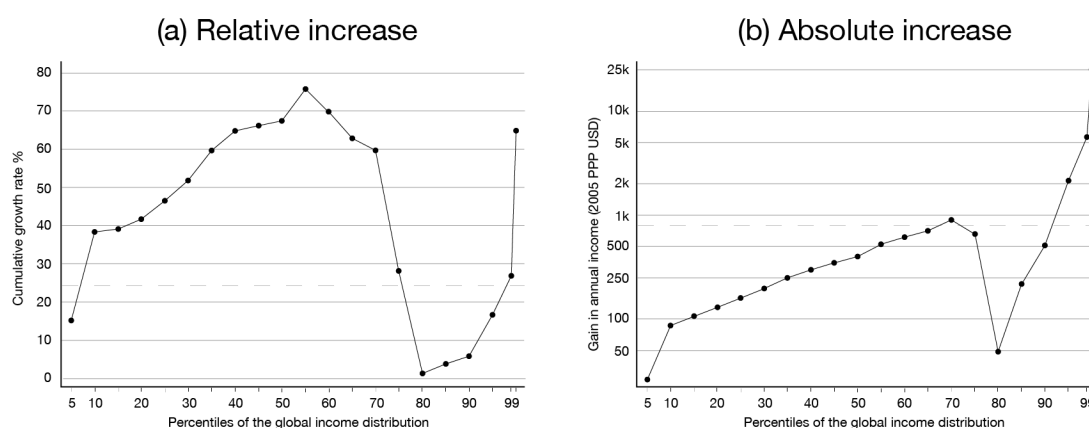
ranks of income distribution, countries, and data surveys.⁷ Lakner and Milanovic (2016), for instance, check the robustness of their estimates of global inequality to different assumptions about the allocation of top incomes missing from survey data. Under the most extreme scenario does the global Gini fail to fall throughout 1988-2008. Under more plausible scenarios a fall is still estimated, albeit more limited.

On top of these considerations it must be born in mind that the top 1% share *as an metric* – regardless of the welfare concept or quality of the underlying data – captures movements across the distribution very differently to the Gini. Namely, it is indifferent to redistribution occurring below the threshold. A country with very high levels of poverty alongside a wealthy middle class will be judged to be just as equal as a second country in which the gap between the poorest and the middle is much narrower – provided only that those above the threshold share similar riches proportionate to the rest. The Gini coefficient, on the other hand is sensitive to relative movements along the entire distribution. Some argue that it is *overly* sensitive to changes in the middle of the distribution relative to those at the extremes, thereby further understating the effect of rising top incomes (WIR 2018).

What is clear though is that the choice of metric reflects fundamentally normative considerations (Atkinson, 1970). Ravallion (2017) points out that, even without any correction for missing top incomes, by using an Atkinson index with a sufficiently high ‘inequality aversion parameter’, the direction of travel of global inequality may be reversed.

Similar considerations hold for the fact of our having based the preceding discussion on *relative* income differences. But it is by no means clear that this is what matters to people universally: *absolute* differences may hold more salience for many, perhaps particularly in a global context (Atkinson and Brandolini, 2010). Fig. 4 presents the cumulative increase in income of various fractiles of the global distribution between 1988 and 2008, taken from Lakner and Milanovic (2016).⁸ Panel (a) presents the figures in terms of percentage increase. Panel (b) shows what this meant in terms of an absolute increase in annual income.

Fig 4: Relative and absolute gains along the global distribution, 1988–2008



Source: Reformatted from Lakner and Milanovic 2016


⁷ The careful works by Burkhauser et al. (2012) and Burkhauser et al. (2018) for the US and the UK, respectively, constitute two notable exceptions.

⁸ Based on broadly the same survey data as used in Figs. 1 and 2

Much has been made of the success experienced by the ‘global middle class’, as captured by the dramatic-looking hump around the 50th percentile. However, the absolute increases for those in this range – around \$300-500 per year – may have seemed less impressive to the individuals concerned had they been aware that those in the top percentile, whilst experiencing similar *relative* growth, saw their annual incomes rise by \$25,000. By a relative measure, the two groups moved in tandem. Anyone sensitive to absolute differences might feel uncomfortable with that characterisation however. By extension, inequality metrics that respond to absolute differences, such as the Absolute Gini, unsurprisingly show continuous rises in global inequality (Niño-Zarazúa et al, 2017).

Conversely, such small absolute increases have also brought about an historic reduction in extreme poverty. Fig. 3 shows the global distribution of income (reconstructed from survey data as before) with the international poverty line indicated.⁹ We see a clear shift from a bimodal towards a unimodal distribution, as incomes in Asia, and China in particular, have increased in relative terms. This is precisely the convergence picked up by the falling global Gini.

Fig 5: The global income distribution, 1988 and 2011




Source: Lakner and Milanovic 2016.

⁹ Note that Fig. 5 is shown in 2005 PPPs. As such the poverty line selected is the \$1.25 level in the same prices in use before the update to \$1.90 in 2011 PPPs.

Note: The distribution is constructed from national household survey data, and is expressed in 2005 PPPs. The red line marks the \$1.25 poverty line in 2005 prices, adopted as the international poverty line prior to the \$1.90 level in 2011 PPPs.

Focussing-in on the shift at the bottom end of the global distribution, Fig. 6 shows how this translated into the number of people in extreme poverty. Between 1981 and 2013 the number living on less than \$1.90 a day (in 2011 international prices) more than halved and the number living on less than \$1.25 reduced by three-quarters. The reduction over this period has been driven by economic growth, despite rising within-country inequality. However, as within-country inequalities become more prominent in the global income distribution, going on to meet Sustainable Development Goal 1 of ending extreme poverty by 2030 will be highly dependent on future trajectories in within-country inequality. Projections by Lakner et al. (2014) suggest that meeting the target will only be possible with continued high growth rates within an inclusive scenario in which the incomes of the bottom 40% grow at 2 percentage points faster the mean in each country.

Fig 6:



5. From global to local trends

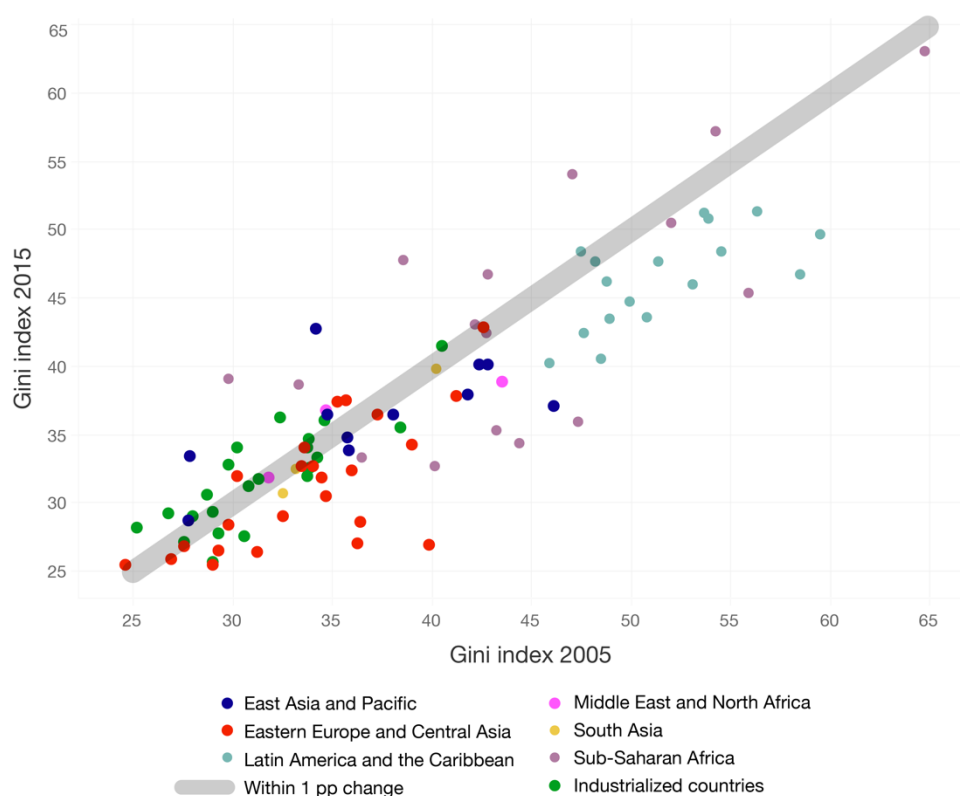
It is important to see that the movements in the global average Gini shown in fig. 1 masks considerable variation across countries and over time. We can get a sense of this heterogeneity by looking at a tabulation of where the Gini was rising or falling by more than

1 point across different world regions in the decade up to 2015 within the POVCAL database, as presented in table 1 and fig. 7.

Table 1: Change in Gini index by world region, 2005–2015

	Change in Gini 2005–2015			
	Number of countries			
	+ >1pp	+/-	- >1pp	Total
East Asia and Pacific	3	1	7	11
Eastern Europe and Central Asia	3	6	15	24
Latin America and the Caribbean	0	2	15	17
Middle East and North Africa	1	1	1	3
South Asia	0	2	1	3
Sub-Saharan Africa	6	2	8	16
Industrialized countries	7	8	7	22
World	20	22	54	96

Fig 7: Change in Gini index by world region, 2005–2015




Source: PovcalNet, April 2018 version.

Note: The table and chart compare the Gini index of surveys falling within two years of the reference year. The surveys refer to both income and consumption, but the difference is always calculated on the basis of surveys of the same welfare concept. Both China and Indonesia appear twice, with rural and urban figures presented separately.

We see a much greater number of countries have lower Gini coefficients around 2015 than around 2005, in line with the falling average Gini seen earlier. In no world region do we find more countries with an increasing Gini than a decreasing one. There are, nonetheless clear differences across world regions. Industrialized countries were evenly split between increasing, decreasing and stable Ginis according to POVCAL estimates. In contrast, almost every Latin American country saw significant declines in Gini index over this period as charted in fig. 8 for a selection of countries, based on SEDLAC figures for disposable household income.

Fig. 8: Gini Index in Latin American countries

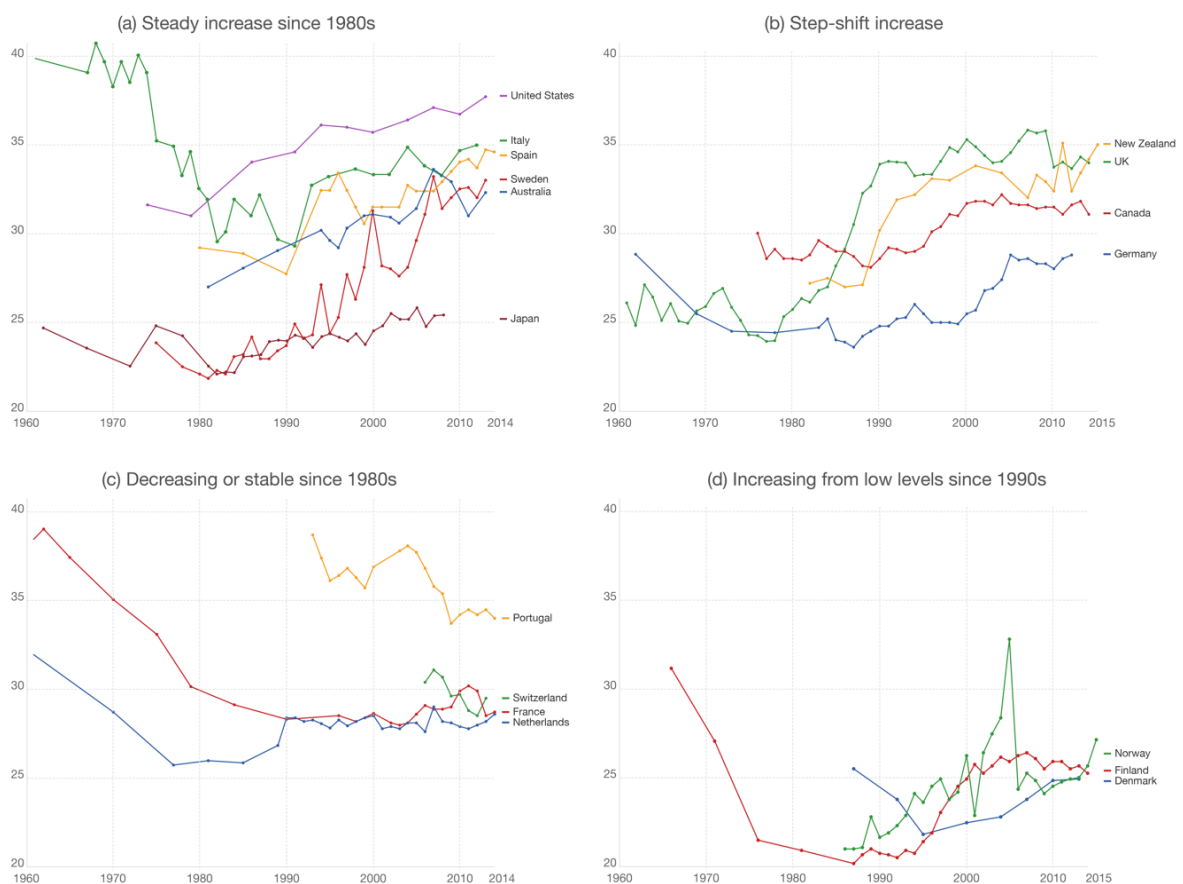


Source: SEDLAC

In the case of high income countries, the availability of longer-run data allows us to make this point about heterogeneous pathways all the more clearly. The general rise in inequality from the 1980s seen in many high income countries is well known. But there are important differences even here in terms of the timing and extent of any increase. Fig. 9 groups high income countries into those following broadly similar trends over this period based on Gini coefficient of household disposable income. A number of countries, including the US, underwent a more or less continuous increase in inequality between the 1980s and the 2010s (Panel A). Beginning their climb somewhat later in the 1990s, and from a lower starting point, several Nordic countries form another cluster (Panel D).

Indeed a geographic distinction is often made between high-inequality Anglophone countries, the more moderate Continental European countries and low-inequality Nordic countries. Whilst informative, even this very loose typology masks important differences. Sweden, for instance (Panel A), stands out in the extent of the rise in inequality seen there, pulling away from its Nordic neighbours and into the ranks of countries like the UK and New Zealand. More modest, but still significant increases in Finland, Denmark and Norway stand in contrast with the steady levels in France and the Netherlands (Panel C) which contribute to a relative convergence between Continental Europe and Nordic countries since the 1980s. The step increase seen at higher levels in the UK, Canada and New Zealand and at lower levels in Germany and Finland (Panel B) also merits attention; and the levelling-out in these countries increasingly serves to accentuate the exceptionalism of the US among rich nations.

Fig. 9: Different trajectories of Gini index in high-income countries



Source: *Chartbook of Economic Inequality* (2017); data for US and Denmark is from *LIS Key Figures*
Notes: In most cases figures refer to disposable (after taxes and transfers) household income, equivalized for household composition. For Canada, the unit of analysis in the family. For Italy, figures are per capita.

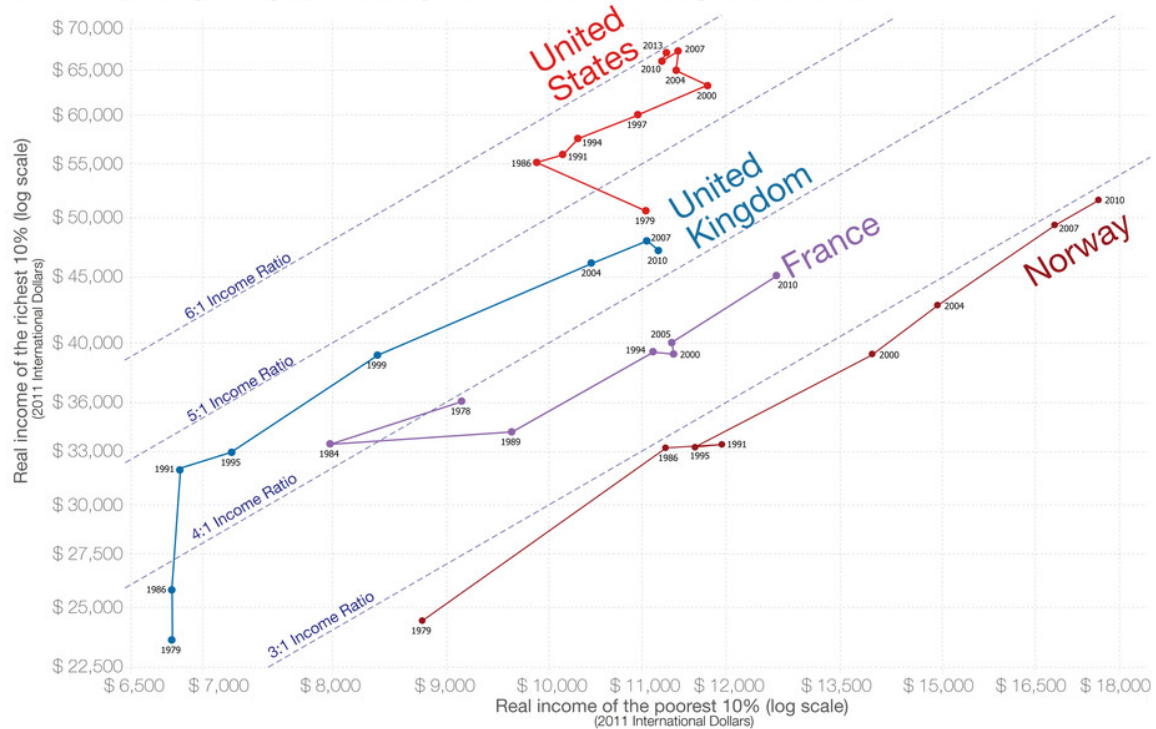
Fig. 10 makes a similar point, but considers the absolute incomes of the bottom and top decile of the US, the UK, France and Norway. The chart is useful in being able to show the effect of both economic growth and relative changes in the top and bottom deciles on the absolute incomes of both groups. Again we see that, even among these relatively similar

countries, exposed to global market forces in very similar ways, the extent to which they have been able to promote inclusive growth differs remarkably.

Fig 10:

Income growth of the poorest 10% vs income growth of the richest 10%

Incomes are real disposable household incomes. Shown is the income cutoff between the richest and poorest 10% and the rest of the population. Incomes are adjusted for price changes over time and for price differences between countries (expressed in international dollars).



Data source: 'Incomes across the Distribution Database' by Stefan Thewissen, Brian Nolan, and Max Roser. Based on LIS data. The data visualization is available at OurWorldinData.org. There you find the raw data and more visualizations on inequality and growth. Licensed under CC-BY-SA by the author Max Roser.

The US again stands out in terms of high levels of inequality, with the ratio of the top and bottom deciles increasing from 4.5 to 6 since 1979. But we can also see how this has impacted on the absolute incomes of the bottom decile which have remained stagnant for over 30 years. In the opposite corner we find Norway, in which the incomes of the poorest 10% have tracked those of the richest 10% in relative terms through strong growth, such the incomes of both groups doubled over the same period. We see that, whilst the top decile in the US is significantly richer than in Norway, the bottom decile is considerably poorer. Similar conclusions hold for France, though as compared to Norway, income trends move along at a higher 90:10 ratio and the incomes of both groups have grown somewhat less. The UK shows a marked increase in inequality through the 1980s, with a decade-long freeze on the incomes of the poorest 10%, though beginning from much lower levels of inequality and much lower incomes than the US. Since the early 1990s however, the incomes of both groups have grown considerably, with inequality falling moderately since the millennium according to this measure.

6. Conclusion

The aim of this paper has been to set recent changes in inequality within countries in a long-run global perspective. By pooling national household surveys across the globe we have seen that the average Gini rose through the 1990s but has either levelled out or fell according to whether we look to the weighted or unweighted mean. Aggregating over such a broad set of data, raises important data quality issues, in addition to those already present at the level of individual household surveys. In particular, I have flagged as a particular worry in lack of comparability between consumption and income surveys, as well the problem of 'missing' top incomes, which is almost certainly biasing estimates downward, and possibly understating recent trends too. In any case, however, the variation in within-country inequality across the globe has been modest compared to the fall in between-country difference over the same period, at least based on the Gini index. Indeed, a more radical and resounding counterpoint to this benign depiction of inequality trends may be found by adopting an alternative measure of inequality: both the top 1% shares in pre-tax income and the absolute Gini show protracted rises over the same period. Judged solely on the basis of extreme absolute poverty, however, those same shifts in the global distribution represent an extraordinary and historic movement in right direction.

it is in individual countries, though, that inequality makes itself felt and policies to reduce it are enacted. Whilst globalisation has no doubt entwined the fate of nations' fortunes together to a far greater degree than ever before, the heterogeneity visible in the levels and trends of inequality, across countries and as represented by different measures, demonstrates the large degree of autonomy that individual countries continue to possess in this domain. Whilst a global perspective on inequality certainly provides an interesting context, this suggests that much more may be learnt about how national policymakers may respond to global economic forces from an assiduous attention to the changing details along the income distribution within different individual countries.

References

Alstadsæter, Johannesen, Zucman 2017. Tax Evasion and Inequality. NBER working paper 23772.

Atkinson, A. B. 1970. "On the Measurement of Inequality". *Journal of Economic Theory* 2: 244-263

Atkinson, A. B, Bourguignon, F, 2015. Introduction to *Handbook of Income Distribution, Vol 2*, Atkinson and Bourguignon (eds.). Handbooks in Economics. Amsterdam: North-Holland.

Bourguignon, F. (2016). *The Globalization of inequality*. Princeton University Press.

Atkinson, A. B., Brandolini, A. 2010. "On Analyzing the World Distribution of Income." *The World Bank Economic Review* 24 (1), 1–37

Lakner, C., et al. (2014) 'The role of inclusive growth in ending extreme poverty', http://www.ecineq.org/ecineq_lux15/FILESx2015/CR2/p191.pdf

Lakner, C., Milanovic, B., 2016. Global Income Distribution: From the Fall of the Berlin Wall to the Great Recession. *World Bank Economic Review*. 30 (2): 203–232

Milanovic, B., 2014. "All the Ginis, 1950–2012 (up-dated in Autumn 2014)." November, World Bank, Washington, DC. <http://go.worldbank.org/9VCQW66LA0>.

Morelli, S., Smeeding, T., Thompson, J. 2014. "Post-1970 Trends in Within-Country Inequality and Poverty: Rich and Middle Income Countries", in *Handbook of Income Distribution, Vol 2A*, Atkinson and Bourguignon (eds.). Handbooks in Economics. Amsterdam: North-Holland.

Niño-Zarazúa, M., Roope, L., Tarp, F. 2017. Global Inequality: Relatively Lower, Absolutely Higher. *Review of Income and Wealth*. 63 (4): 661-684

World Bank 2016. *Poverty and Shared Prosperity 2016: Taking on Inequality*

World Income Report 2018. 2017. World Inequality Lab. Available at <http://wir2018.wid.world>