No Elephant in the Room? Noisy Convergence and the Global Growth Incidence Curve*

First version: February 2019. This version: December 2019

WORK IN PROGRESS. Click here for the latest version

Lucas Herrenbrueck (Simon Fraser University)

JEL: C18, C22, O47

Keywords: Elephant curve, globalization, war on poverty, convergent growth, quantile analysis, selection bias

Contact: herrenbrueck@sfu.ca

In recent decades, worldwide incomes have grown enormously. But how has this growth been distributed? One way to answer this question is with a "growth incidence curve" (GIC, henceforth): percentiles of the income distribution are shown on the X-axis, and the percent change associated with each percentile on the Y-axis. Applied to worldwide data collected for the period 1988-2008, the percentiles in the middle and at the top of the world income distribution have grown the fastest, while percentiles at the bottom and around 80 have stagnated [7, 8] (see Figure 1). This striking pattern in the global GIC was dubbed the "Elephant curve" and, in media coverage, was commonly interpreted as follows: globalization has *redistributed* incomes from all but the richest people in rich countries towards the fast-growing middle classes in emerging economies, especially China, and this pattern of global redistribution may explain populist resentment in rich countries.

However, I propose that there may be a simpler, statistical, explanation for the Elephant pattern. Consider a random ("noisy") income growth process affecting individuals whose initial incomes follow a bimodal distribution. Random growth smoothes out the income distribution, which explains the hump of the Elephant curve: given that the 70th percentile was the trough in the bimodal world income

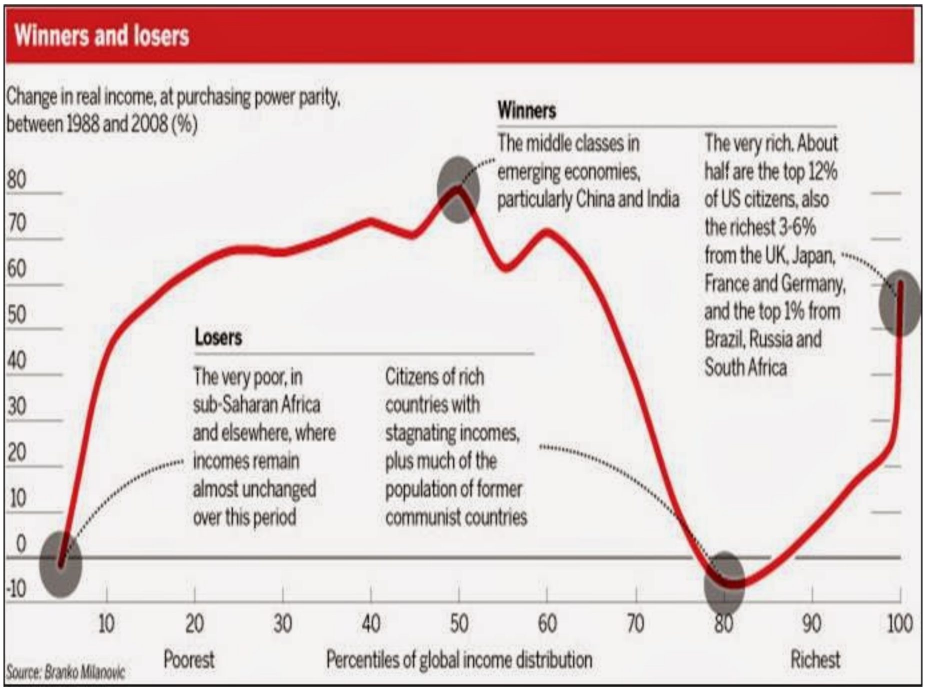

Figure 1: The original Elephant curve, as presented in the media. Data is based on Lakner and Milanovic [7], figure and interpretation from the Financial Times (2013).

distribution of the 1980s (see Figure 3), people just below it in 2008 are positively selected for growth, while people just above it in 2008 are negatively selected. However, random growth alone would predict that the poorest quantiles would experience negative growth. Explaining the fact that they did not – they experienced zero growth in the original version of the Elephant curve, and robustly positive growth in more recent studies that use better data (see Figure 2) – requires a modest amount of convergence, whereby average income growth is highest for people near the bottom of the global distribution and lowest for people near the top. I simulate a growth incidence curve with these three ingredients – a bimodal distribution as the starting point, a little bit of convergence, and a lot of randomness – and I find that it almost exactly replicates the one observed in the macro-data.

The implication of this alternative model is that the Elephant curve may just be an artifact of noisy growth during a historically unique era. Contra most of the discussion surrounding it, the Elephant curve therefore does *not* constitute evidence that incomes were redistributed in a quantile-biased way, where the 60th percentile 'took' from the 80th percentile, and fast growth in emerging markets (prominently, China) came 'at the expense of' the lower and middle classes in developed countries.

For a Western observer, it is easy to notice China's growth but forget about selection. China, India, and the Congo (Dem. Rep.) started out with similar incomes in 1988; China grew fast, India more slowly, and war-torn Congo collapsed. Randomness means that *someone* poor in 1988 was going to hit on the right combination of circumstances and policies for fast growth; it happened to be China, but in 1988, we did not know it would be China.

Absent world-wide micro-data tracking individual families through the decades, we cannot conclusively prove whether the quantile-biased-growth hypothesis or the statistical-artifact hypothesis is correct. But the latter is more parsimonious, and thus deserves to be studied more carefully. In particular, one of its three ingredients is that individual growth experiences are subject to a large amount of randomness, much of which may be between-countries but some of which must be within-country. Existing micro-data could be used to quantify these sources of randomness.

What about growth at the top of the income distribution – how can noisy convergence explain the raised trunk of the Elephant? First, it is important to note that the data underlying the original Elephant curve [7] stops at the global 99th percentile, thus excluding the top one percent. But these are the *global* one percent, a much less exclusive club than the U.S. One Percenters that have attracted so much attention; the 99th percentile of the world income distribution only corresponds to about the 88th percentile of the U.S. income distribution [6].[1] Given that incomes at the 88th percentile in rich countries are not known to have grown especially fast on average, this is suspicious.

The answer is again selection bias. For someone to be part of the worldwide 99th percentile, they can take one of two paths: (a) start above the 99th percentile and experience low growth; (b) start poor and strike it rich. But (a) applies to few people, while (b) applies to many. Thus, the growth of the *new* 99th percentile is largely determined by group (b). This implies that the raised trunk of the elephant is also a statistical artifact, one which tells us little about the average growth experience of high income earners, and which should be seen as irrelevant to the discussion about income shares of top earners (as in, the U.S. top one percent or top 0.1 percent).[2]

[1] Most (Western) economics professors are part of the global one percent, and thus too rich to be represented in the original Elephant chart!

[2] To be clear, an increasing income share of top earners (for which there is separate evidence [1]) is sufficient to cause a steep right end of the GIC. But it is not necessary – random noise has the same effect. And rising income shares of the U.S. top one percent are mostly driven by rising income shares of even more exclusive groups, such as the top 0.1 or 0.01 percent [12].

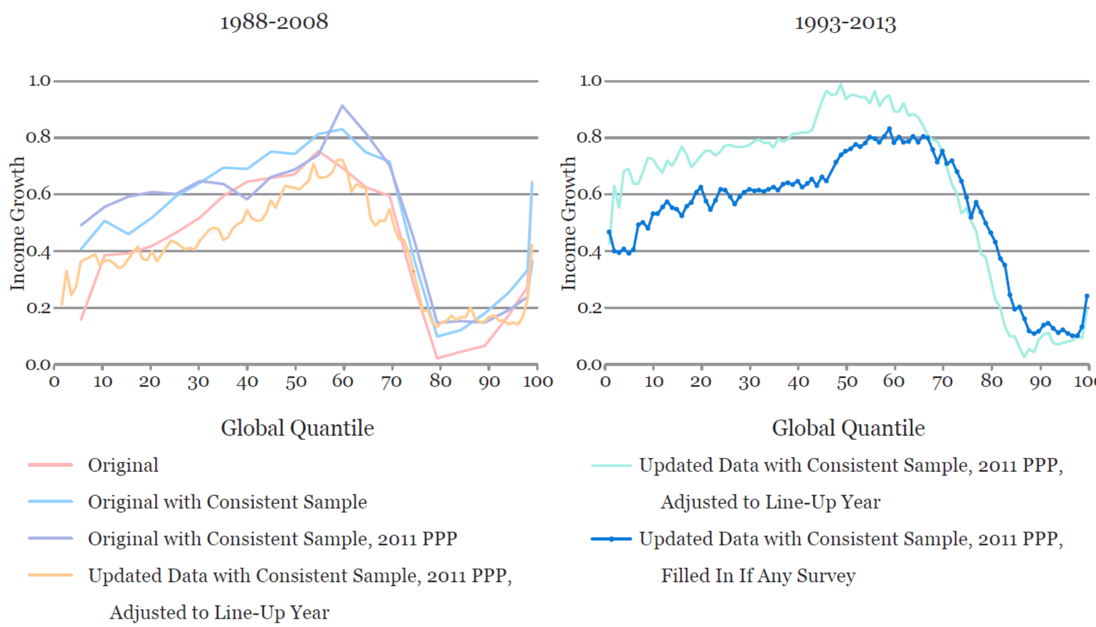

Figure 2: Slicing and dicing the data; details change, but the Elephant pattern persists. Figure taken from Kharas and Seidel [6].

My simulation also has implications for the bottom end of the distribution. Just like the raised trunk of the Elephant reflects that high growth is the likeliest way to get to the top, the lowered tail of the Elephant reflects that low growth is the likeliest way to end up at the bottom. Thus, the fact that growth of the global 5th percentile is fairly low is not evidence against convergence. This fact has been pointed out before and is supported by country-specific studies using household panel data: an upward-sloping GIC near the bottom is consistent with convergence in expectation, given the presence of random noise [5, 4].

The problem from a development perspective is that growth in expectation does not necessarily mean growth for everyone. As a consequence, if the goal is to eliminate global poverty, growth-promoting policies alone will not do the job. Also needed are policies that reduce variance: that is, a global social safety net.

Many authors have noted that growth incidence curves such as the Elephant curve, which plot the percent change of a given quantile of the distribution, give misleading impressions of individual growth experiences when the ranking of individuals is not constant [10, 5, 4]. Since the composition of the various quantile groups is not tracked, such a curve is called "anonymous" in the literature [2]. The alternative is called a "non-anonymous" GIC: we plot the growth of individuals (or groups)

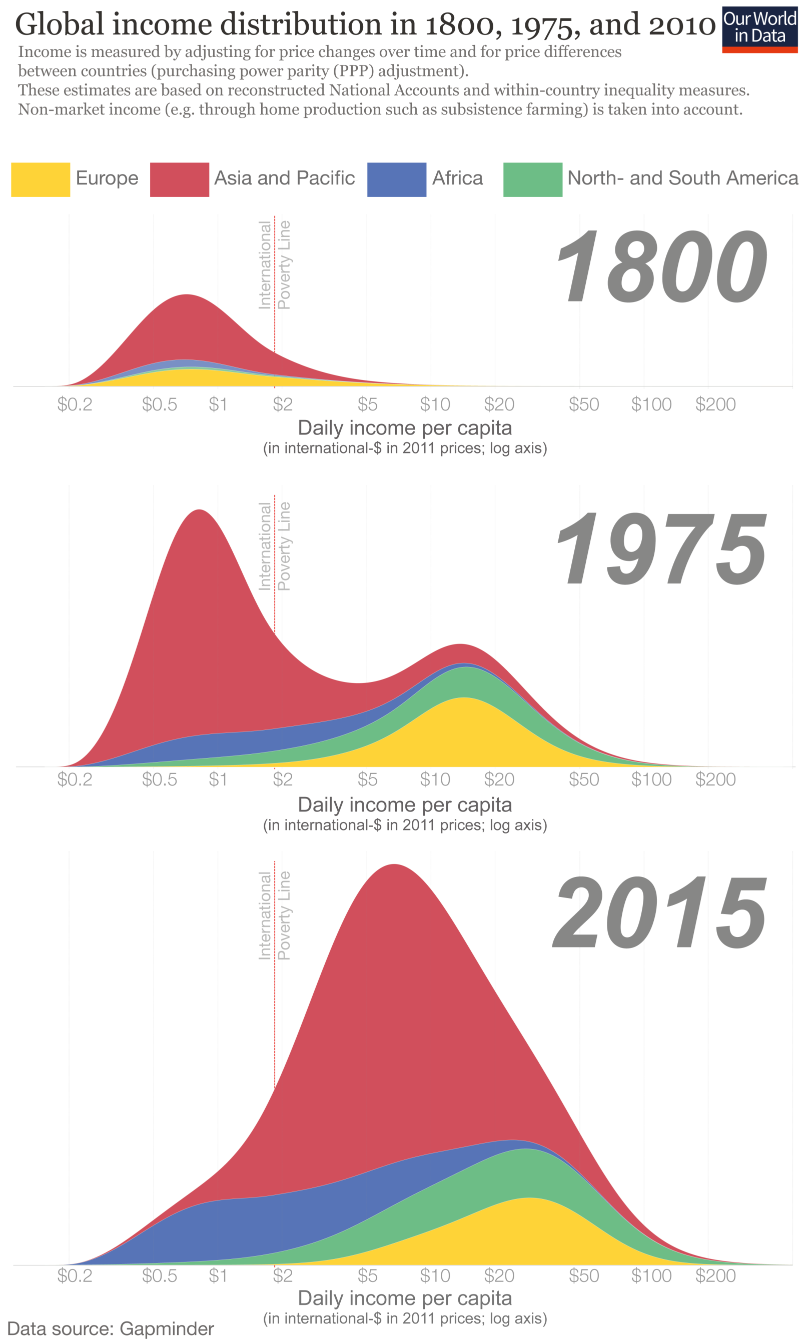

Figure 3: The global income distribution over time. Note how unusual the late-20th century income distribution was, being bimodal. Countries either belonged to a developed first/second world, or an undeveloped third world. This is no longer the case in the 2010s.

corresponding to an initial quantile of the income distribution, but hold the ranking constant [2]. The problem with historical global income data is that the household panels that would allow us to track individual income groups simply do not exist. Thus, anonymous growth incidence curves may be the best we can do.

This problem is acknowledged in the academic literature on the Elephant curve [7, 8, 3, 9, 6]. However, even sophisticated discussions sometimes muddle the distinction (e.g., [11]: "the biggest losers. . . of globalization were those between the 75th and 90th percentiles of the global income distribution whose real income gains were essentially nil"), and it is completely ignored in media coverage ("the winners and losers from globalization"; "this chart explains Trump and Brexit" [[cites]]). While attempts can be made to construct a "quasi-non-anonymous" GIC by holding country-income groups constant (rather than individuals, which we would ideally want) [8, 6], the results are fragile to assumptions and at any rate incomplete.

To my knowledge, this is the first paper attempting to *reverse-engineer* the individual-level growth process that would be consistent with a given GIC. To keep the exercise disciplined, I only tune three parameters: the level and slope of the expected growth rate (a negative slope implying convergence in expectation), plus the standard deviation of individual growth outcomes. A positive expected growth rate is necessary to match positive growth on average, a positive standard deviation is necessary to generate ranking reversals and thus the Elephant pattern, and a force towards convergence is necessary to explain why growth at the bottom is still positive despite negative selection. As I will show, an Elephant pattern emerges in the anonymous GIC which closely resembles that in the data, even though 'true' expected growth is monotonically decreasing in initial income.

Finally, I explore robustness of the simulation by simulating the same growth process with three alternative initial income distributions. When the initial income distribution is normal, the anonymous GIC is nearly flat. When the initial distribution is uniform, the Elephant pattern appears; this implies that the Elephant pattern does not require the initial distribution to be bimodal per se, just "thinner around the middle" than a normal distribution. Third, I take the *outcome* of the original simulation as the starting point of another simulation. Thus, this can be interpreted as a forecast of income growth over the next generation. The Elephant pattern persists but is weaker than in the original simulation, which is consistent with data collected after 2008 (see Figure 2 and [6]).

1 The model

Time $t = \{0, \dots, T\}$ is discrete. There is a large population of individuals $i = \{1, \dots, N\}$, each with income $Y_{i,t} > 0$ at time t. Define $y_{i,t} \equiv \log(Y_{i,t})$ so that differences in y correspond to growth rates of Y.

For the initial distribution, draw N values of $y_{i,0}$ randomly from a distribution with cumulative distribution function $F(y)$. Each period, individual incomes evolve as follows:

$$y_{i,t+1} = y_{i,t} + g(y_{i,t}) + e_{i,t} \tag{1}$$

where $e_{i,t}$ are i.i.d. draws from a normal distribution with standard deviation $\sigma > 0$, and $g(y)$ is a growth function that specifies the expected growth rate for a particular level of income. (That is, growth depends on the level of income but not on one's rank in the income distribution.) Here, I use:

$$g(y) = \max\{0, \gamma - \delta y\} \tag{2}$$

I use this functional form because it satisfies: (a) positive growth on average if $\gamma > 0$; (b) convergence in expectation if $\delta > 0$; (c) non-negative expected growth for everyone. Other functional forms could be chosen that satisfy these criteria, and it is not obvious that criterion (c) is even necessary for the conclusions, but at the level of analysis employed at this point it does not appear that more precision is necessary.

For the simulations in this paper, I use $N = 100{,}000$, $T = 20$, $\gamma = 0.06$, $\delta = 0.015$, and $\sigma = 0.335$. For the initial income distribution $F(y)$, I use a mixture distribution between F_1 and F_2, where $F_1\,\mathcal{N}(-1/0.7, 1)$ and $F_2\,\mathcal{N}(1/0.3, 1)$, F_1 has weight 0.7, and F_2 has weight 0.3. For interpretation, this means that:[3]

- Mean log income is 0, and expected growth at the mean is 6 percent per year
- Median log income is -0.86, and expected growth at the median is 7.3 percent per year
- Expected growth is zero for the 93rd percentile and above
- The standard deviation of log income is 0.335 over one year, and 1.5 over 20 years

Figure 4 shows the initial income distribution and the expected growth process.

[3] Both the mean and standard deviation seem high if the period is a "year"; future versions of this paper will refine this. If the reader prefers, they could be though of as representing growth over 1.5 or 2 years instead.

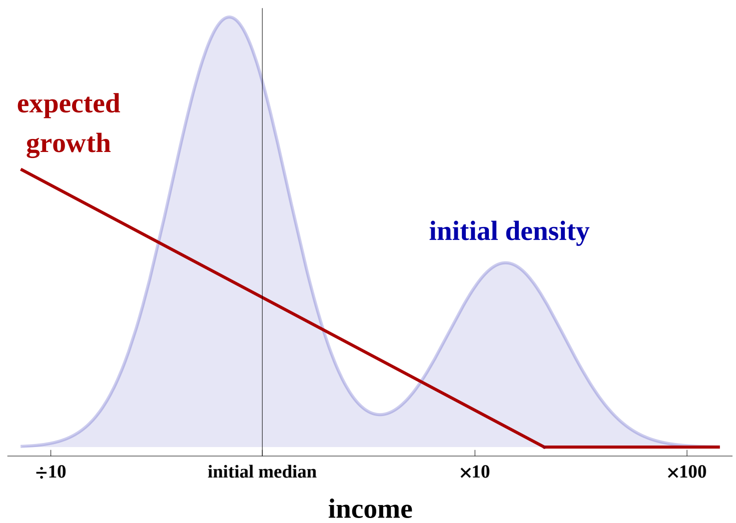

Figure 4: Assumptions of the model: initial (log) income distribution, and expected growth rate over the next year.

2 Simulation of global income growth

Figure 5 shows the simulated densities of the initial income distribution and the income distribution after 20 years of noisy-convergent growth. This looks as expected: the two peaks of the initial distribution have merged into a single one (reflecting noise), and the overall distribution has shifted slightly to the right (reflecting growth). Nothing uncanny is going on here. So in order to understand how this can produce an Elephant pattern in the GIC, we need to look at quantiles and at individual outcomes.

Figure 6 does exactly that. Each dot represents one of $N = 100{,}000$ simulated individuals, and the vertical axis shows that individual's growth outcome. The only difference between the figures is the X-axis. In the left panel, individuals are ranked by *initial* income, so that the best-fit trendline can be interpreted as the average growth outcome of an individual starting at a particular point in the income distribution. Not surprisingly, this trendline is downward sloping, looking like a smoothed version of the expected growth curve in Figure 4. In the right panel, by contrast, individuals are ranked by *final* income. Now, the best-fit trendline is mostly upward sloping: individuals who end up at the top most likely experienced fast growth, while individuals at the bottom most likely experienced negative growth.

But the most interesting shape is in the middle, where the scatterplot of individu-

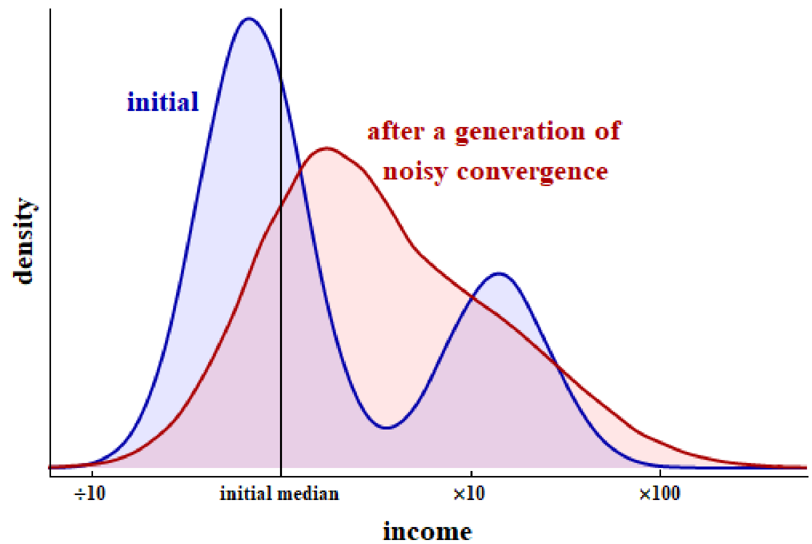

Figure 5: Simulated income distributions: initial, and after 20 years.

als separates into two distinct clouds. The upper-left one represents the "poor" part of the initial income distribution, while the bottom-right one represents the "rich" part. Clearly, the difference in average growth between the two clouds is negligible. But because individuals near the 60th percentile (of the final distribution!) are most likely drawn from the "poor" part, their average growth looks strong, while the opposite is true for individuals who end up near the 80th percentile. This is the source of the Elephant pattern which will appear in the anonymous GIC, because the anonymous GIC scrambles people's initial and final incomes together.

Thus, Figure 7 shows the main results of the paper, contrasting the anonymous GIC (A-GIC, henceforth) with the "expected" GIC (E-GIC, henceforth; expected growth by initial percentile). The A-GIC is constructed in the same way as Figures 4-2, and the E-GIC is just the trendline taken from the left panel of Figure 6.

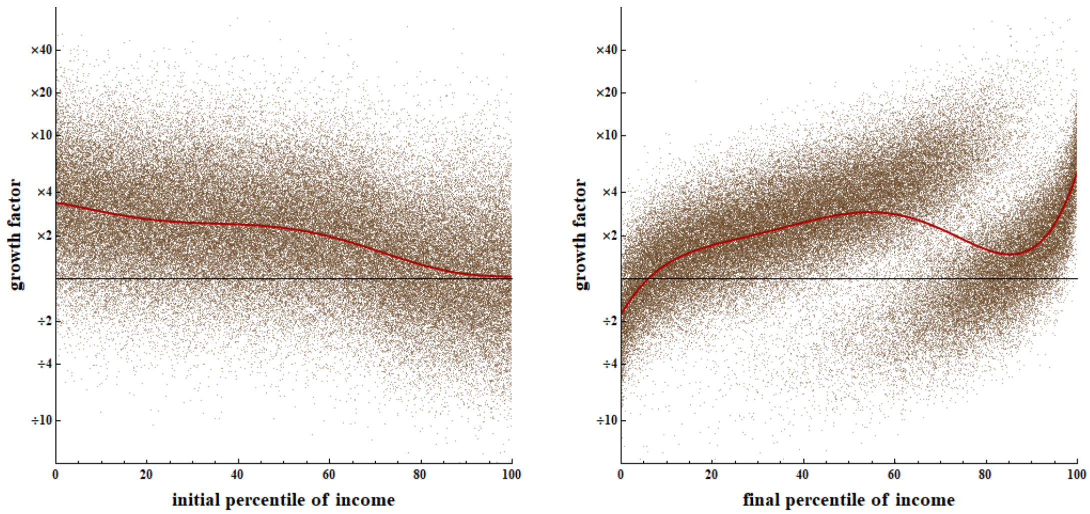

Figure 6: Simulated individual outcomes, by initial percentile and by final percentile

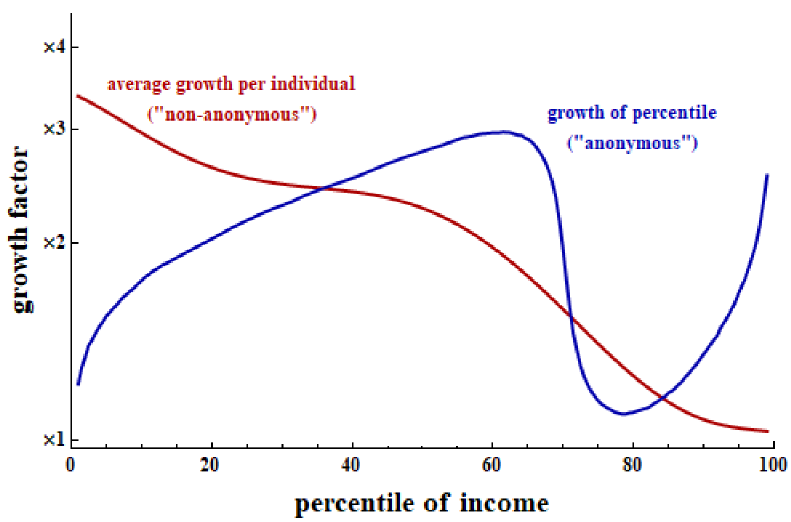

Figure 7: Anonymous versus expected growth incidence curves in the simulation.

3 Implications

Audiences (both media consumers and academics) seem to intuitively interpret A-GICs as E-GICs, but clearly they can look very different. The simulated A-GIC suggests (a) an absence of convergence, since the low percentiles grow more slowly than the median and top percentiles; and (b) redistribution around the 70th percentile. The simulated E-GIC reveals the opposite: (a) the poor grew the fastest on average, and (b) expected growth is monotonically decreasing throughout.

And as Figure 6 further shows, individual growth experiences can be widely dispersed, which makes it easy to cherry-pick examples to support any story. Yes, global poverty persists, but not because poor countries have no hope; it persists because it is continually being 'restocked' by disasters such as wars and famines (consider Haiti, South Sudan, Syria, and Yemen in recent years). Yes, China as a whole grew quickly; but India and Vietnam grew more slowly (from the same starting point), and Latin America fell behind (from a higher starting point). Yes, many individuals from developing countries have joined the global super-rich elite; but if this has come at anyone's expense, it is just as likely the poor in their own countries rather than the U.S. middle class. Or, it may not have come at anyone's expense, but simply be the luck of the draw; some outstanding draws are to be expected when millions of people get to make draws.

3.1 Growth at the very top

Zooming in on the top (Figure 8): clearly, there is no slowdown and the A-GIC just becomes steeper and steeper. The reason is selection – to end up at the top, one must have experienced excellent shocks – coupled with the assumption of normally distributed shocks e_{it}. Because the shocks have no upper bound, someone with extreme luck can always be found if only we zoom in enough. With bounded shocks, we would eventually hit a bound but with normal shocks we do not. No matter what the underlying average growth rate near the top is, the A-GIC always asymptotes to infinity!

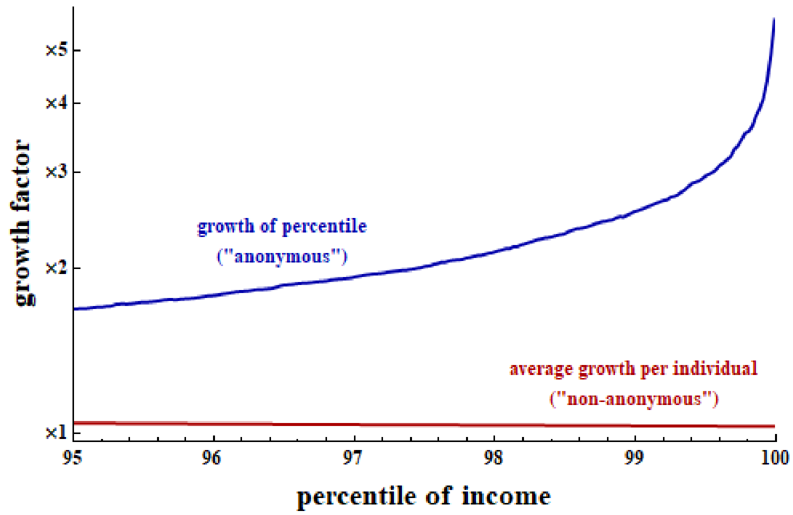

Figure 8: Zooming in on Figure 7 (anonymous versus expected growth incidence curves in the simulation). I drop the bottom 95 percent and fill in percentiles 99-99.9 (notice the expanded Y-axis scale compared to Figure 7).

3.2 Growth at the bottom, and the War on Poverty

What causes poverty: low expected growth, or high variance (specifically, a bad draw from this variance)? It could be both in principle. But as I have argued earlier, low average growth is hard to square with the fact that the left end of the Elephant curve is still positive, despite the issue of negative selection. (Microstudies also support convergent, "pro-poor" growth in various countries [5, 4].)

This points to variance as the explanation. But if variance is indeed the issue, growth-promoting policies alone will never eliminate global poverty. Variance-reducing policies might do it, although care should be taken that they are not implemented in such a way as to reduce average growth. One might consider it unfair to deny poor people the chance to strike it rich, even if this is the price to pay for reducing the chance that they become even poorer. Simply put, what is needed is a global social safety net. See Figure 9 for an illustration.

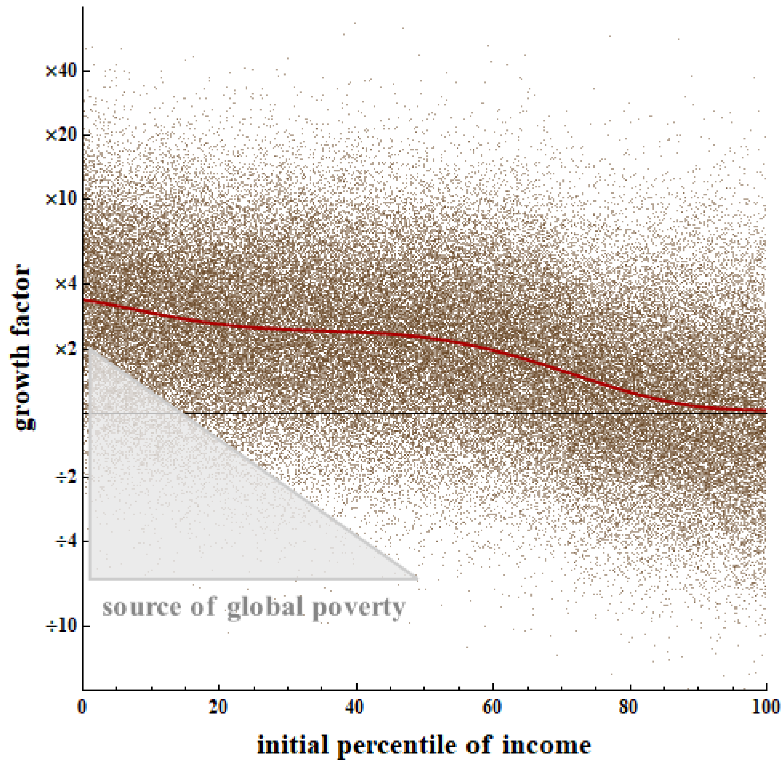

Figure 9: Same as the left panel of Figure 6, but with the (statistical) sources of global poverty highlighted.

4 Alternative starting distributions

4.1 Normal

See Figure 10.

4.2 Uniform

See Figure 11.

4.3 Continuing the main simulation (forecast)

See Figure 12.

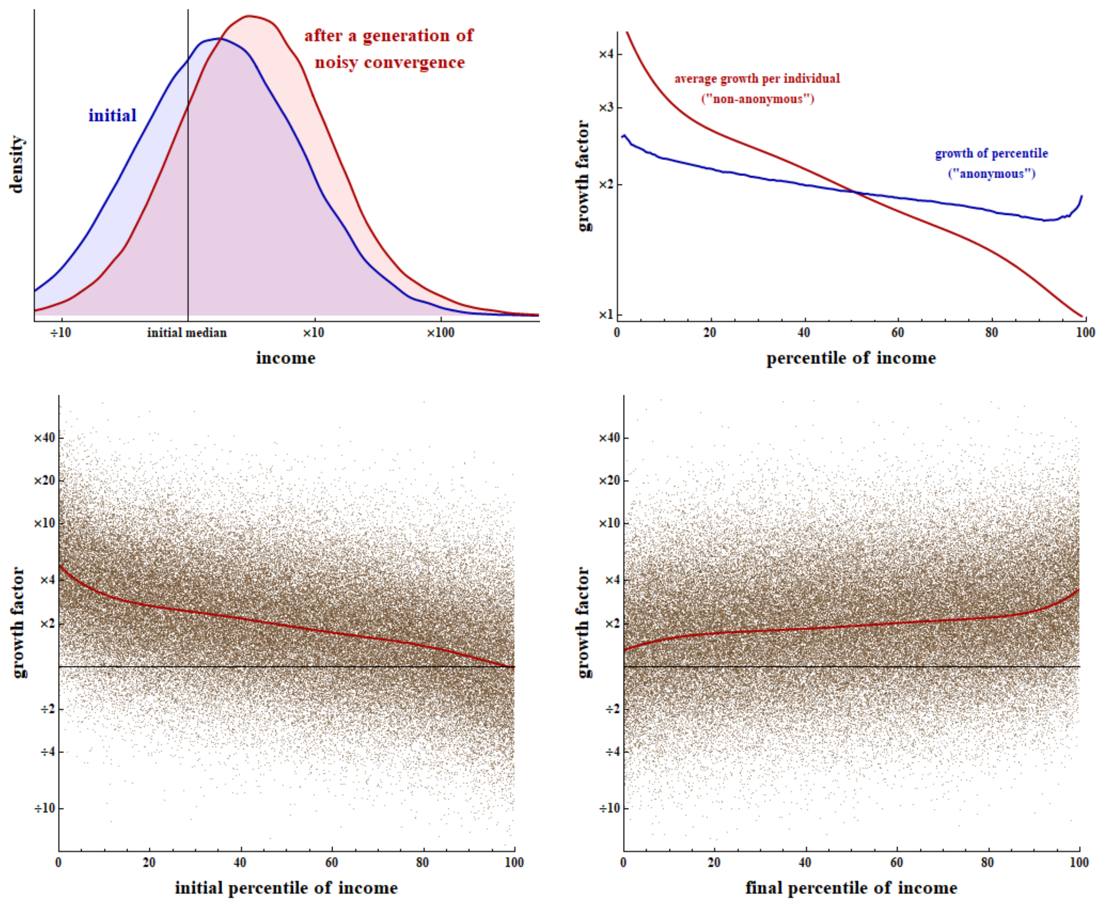

Figure 10: Alternative simulation, with standard normal distribution (over log income) as starting point.

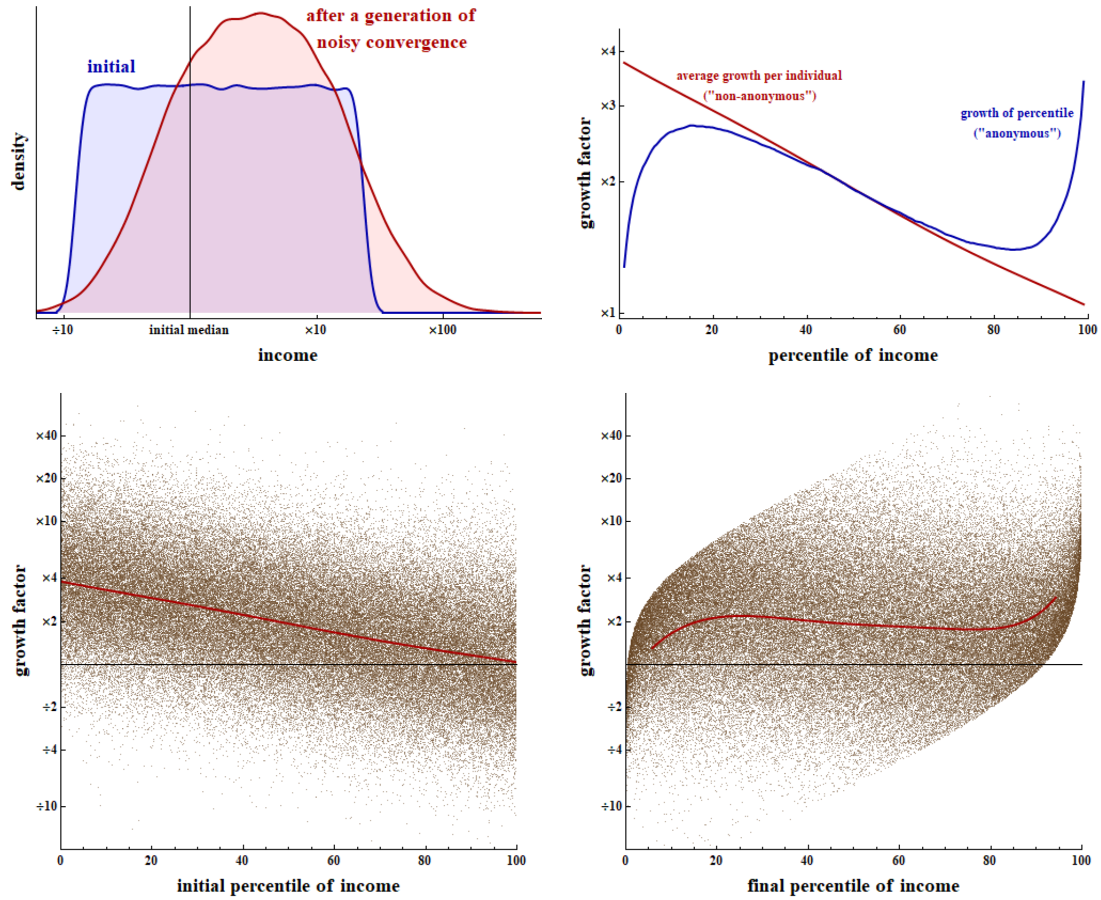

Figure 11: Alternative simulation, with uniform distribution (over log income) as starting point.

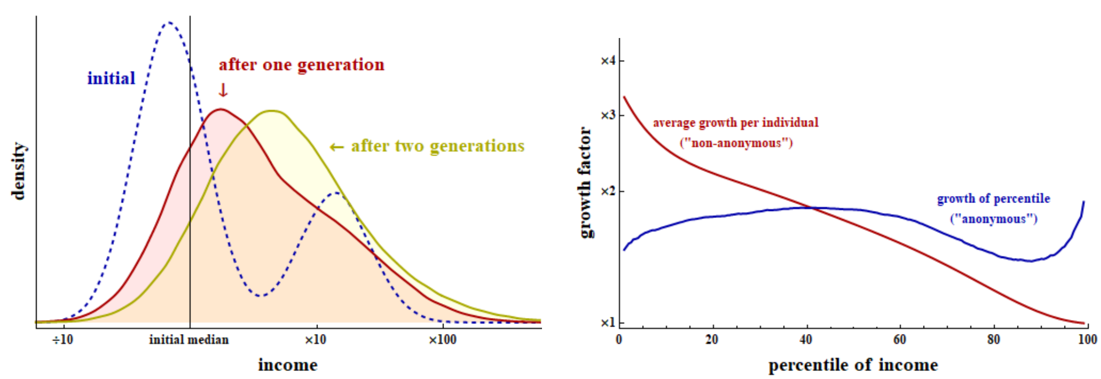

Figure 12: Continuing the main simulation for another generation. Right panel compares generation 2 with generation 1 (not with the initial).

5 Summary

Lesson for development economics: do not trust anonymous GICs unless they can be reverse-engineered by a plausible data-generating process (ideally backed up by micro-data).

Lesson for econometrics: be careful with quantiles. We often think of quantiles (like the median) as 'robust' statistics which are less fragile than means and standard deviations. But as the simulation in this paper shows, this is not always the case. Quantiles like the median are robust to outliers – the funny stuff which happens *on the margins* of a distribution – but not so robust when the funny stuff is happening *on the inside* (like when we are dealing with bimodal distributions).

References

[1] Alvaredo, F., L. Chancel, T. Piketty, E. Saez, and G. Zucman (2018). The elephant curve of global inequality and growth. In *AEA Papers and Proceedings*, Volume 108, pp. 103–108.

[2] Bourguignon, F. (2011). Non-anonymous growth incidence curves, income mobility and social welfare dominance. *The Journal of Economic Inequality 9*(4), 605–627.

[3] Corlett, A. (2016). Examining an elephant: globalisation and the lower middle class of the rich world. *Resolution Foundation Report*.

[4] Grimm, M. (2007). Removing the anonymity axiom in assessing pro-poor growth. *The Journal of Economic Inequality 5*(2), 179–197.

[5] Jenkins, S. P. and P. Van Kerm (2006). Trends in income inequality, pro-poor income growth, and income mobility. *Oxford Economic Papers 58*(3), 531–548.

[6] Kharas, H. and B. Seidel (2018). Whats happening to the world income distribution? the elephant chart revisited. *Global Economy and Development Working Paper* (114).

[7] Lakner, C. and B. Milanovic (2013). *Global income distribution: from the fall of the Berlin Wall to the Great Recession*. The World Bank.

[8] Lakner, C. and B. Milanovic (2016a). Global income distribution: From the fall of the berlin wall to the great recession. *World Bank Economic Review 30*(2), 203–232.

[9] Lakner, C. and B. Milanovic (2016b). Response to adam corletts examining an elephant: globalisation and the lower middle class of the rich world. *CUNY, GC paper, Septiembre*.

[10] Lambert, P. J. (1992). The distribution and redistribution of income. In *Current issues in public sector economics*, pp. 200–226. Springer.

[11] Milanovic, B. (2012). Global income inequality by the numbers: in history and now. Policy research working paper, World Bank.

[12] Piketty, T., E. Saez, and G. Zucman (2017). Distributional national accounts: methods and estimates for the united states. *The Quarterly Journal of Economics 133*(2), 553–609.

[13] Ravallion, M. and S. Chen (2003). Measuring pro-poor growth. *Economics letters 78*(1), 93–99.