VOX

Research-based policy analysis and commentary from leading economists

Search

Columns   **Vox Talks**   **People**   **ePubs**   **Debates**   **Events**   **JMP Vox**   **About**

By Topic   By Date   By Reads   By Tag

Parametric estimations of the world distribution of income

**Xavier Sala-i-Martin, Maxim Pinkovskiy,** 22 January 2010

World poverty is falling. This column presents new estimates of the world's income distribution and suggests that world poverty is disappearing faster than previously thought. From 1970 to 2006, poverty fell by 86% in South Asia, 73% in Latin America, 39% in the Middle East, and 20% in Africa. Barring a catastrophe, there will never be more than a billion people in poverty in the future history of the world.

30   A **A**

World poverty is falling. Between 1970 and 2006, the global poverty rate has been cut by nearly three quarters. The percentage of the world population living on less than $1 a day (in PPP-adjusted 2000 dollars) went from 26.8% in 1970 to 5.4% in 2006 (Figure 1).

Figure 1. World poverty rates
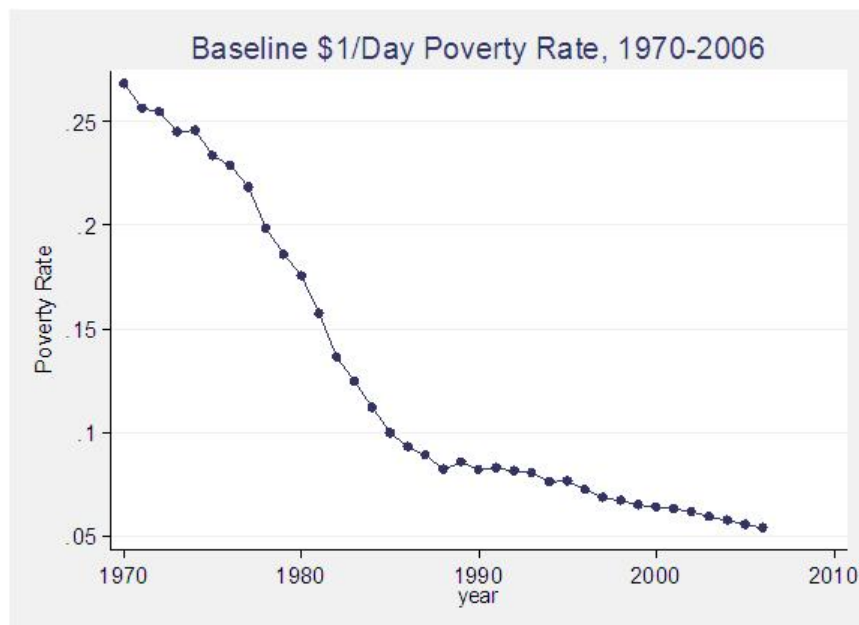
Although world population has increased by about 80% over this time (World Bank 2009), the number of people below the $1 a day poverty line has shrunk by nearly 64%, from 967 million in 1970 to 350 million in 2006. In the past 36 years, there has never been a moment with more than 1 billion people in poverty, and barring a catastrophe, there will never be such a moment in the future history of the world.

New research

In Pinkovskiy and Sala-i-Martin (2009), we compute such numbers by estimating the world distribution of income. It can be estimated by aggregating up all the country income distributions (Figure 2). To compute a country distribution of income, we need data on two elements of this distribution: the average (which is given by GDP), and the dispersion.

Obtaining the average is simple. There exist well-known databases such as the Penn World Tables, the World Bank's World Development Indicators and Maddison's historical statistics that report GDP in real, purchasing-power-parity adjusted terms for almost every country and year.

But obtaining the distribution is more difficult. Data on country income distributions comes from household surveys, conducted by the country, or by international organisations. These surveys take

---

Maxim Pinkovskiy
Department of Economics, MIT

Xavier Sala-i-Martin
Professor of Economics,
Columbia University

Don't Miss

The win-win from eliminating corporate income taxes
Kotlikoff

Why fiscal sustainability matters
Buiter

A consistent trinity for the Eurozone
Buti

Obstacles to international macro policy coordination
Blanchard, Ostry, Ghosh

Regulation, supervision and the role of central banks
Editors

Fed tapering and emerging markets
Eichengreen, Gupta

Global value chains, interdependence, and the future of trade
Lamy

Most Read

This Month | All Time

GDP and life satisfaction: New evidence
Proto, Rustichini

The next sudden stop
Kalemli-Ozcan

Fed tapering and emerging markets
Eichengreen, Gupta

Can we move beyond the Maastricht orthodoxy?
Andor

An economic theory model for amending the Treaty of Lisbon
Basu, Stiglitz

more

Vox Talks

place only for some countries in some years, and they tend to publish very sparse summaries of the data they collect. Fortunately, a compilation of these surveys is provided in the World Income Inequality Database (UNU-WIDER 2008), which is maintained by the UN University, but each survey summarises the income distribution of a country in just five numbers: the quintile shares and the Gini coefficient. Many different distributions of income could produce these five numbers, some of which might imply rising poverty, and others, falling poverty.
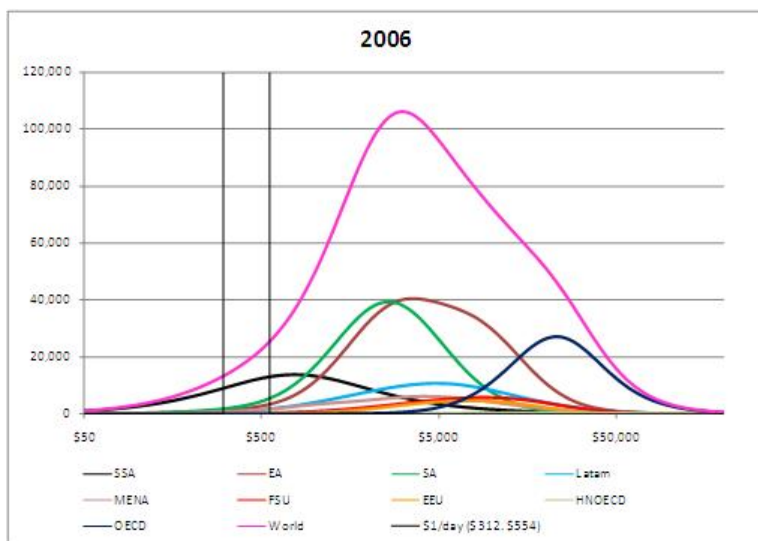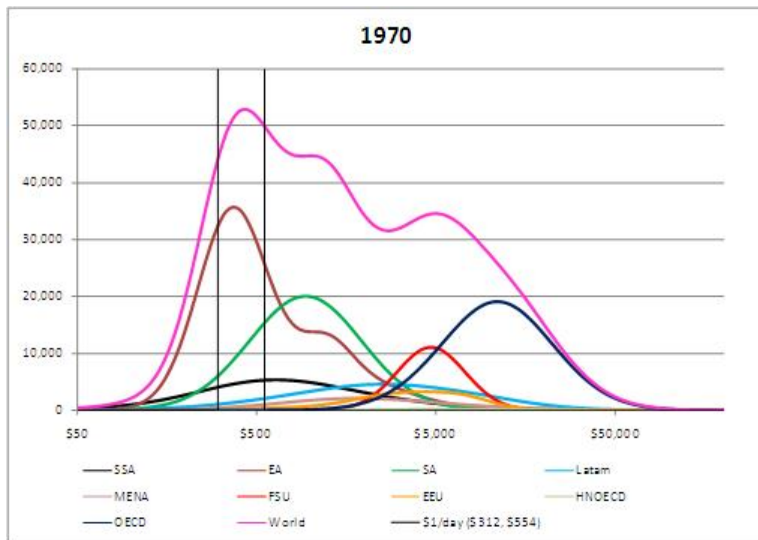
Figure 2. World distribution of income: 1970 and 2006

Our methodology is to recover the distributions of income of all countries by combining the survey data on quintile shares with the assumption that income is distributed lognormally – or that log income has a normal distribution. This assumption has long been employed in the literature on income distributions, at least since Gibrat (1931), who showed that if individuals experience random proportionate shocks to income, the distribution of income will converge to being lognormal in the long run.

We also check how well the lognormal assumption fits our data in comparison with two other functional forms for the income distribution that are used in the literature, the gamma and the Weibull, and find that the lognormal distribution provides superior fit for 98% of the world population.

The lognormal distribution depends on two parameters: the distribution parameter, which completely determines the quintile shares, and the scale parameter, which is used to obtain the mean of the distribution. We estimate the distribution parameter by choosing it to minimise the sum of squared deviations between the quintile shares we observe and the theoretical values of the quintile shares that would be predicted by this parameter, and then, use the distribution parameter and GDP data to recover the scale parameter.

With this parametric assumption, we can recover completely the income distributions of the countries in the years for which inequality data is available. Meanwhile, for those countries where data is not available, we interpolate or impute inequality data, taking advantage of the relative continuity of the inequality series. We then aggregate these distributions to obtain the world distribution of income.

We can compute the poverty rate for any line we like – just by integrating the distribution up to the poverty line.

Differences with the World Bank estimates

We find that in the last 40 years, world poverty rates and counts have collapsed. These results are consistent with Sala-i-Martin (2006), which uses an alternative methodology to recover the world distribution of income, but differ from the World Bank's estimate that 1.4 billion people live in extreme poverty.

These differences exist because:

- Consistent with the Millennium Development Goal to "halve, between 1990 and 2015, the proportion of people whose income is less than $1 a day," we look at income poverty, while the World Bank looks at consumption poverty;
- The World Bank estimates the mean income or consumption from the mean of the household surveys that present the inequality data, whereas we follow the growth literature in estimating the mean income as national accounts GDP per capita.

Our results also show that while the reduction in poverty has been most dramatic in East Asia, the result that global poverty rates decline is not driven by East Asia. Between 1970 and 2006, the South Asian poverty rate fell by 86%, the Latin American poverty rate fell by 73%, the Middle Eastern poverty rate fell by 39%, and even the African poverty rate fell by more than 20%.

We can compute not only the world poverty rates and the poverty rates of any country or region, but also other statistics related to the distribution of income. For instance, we can compute the world gini coefficient, a measure of world inequality, for every year between 1970 and 2006. We show that world inequality measured by the gini fell from 67.6 to 61.2 (Figure 3), and similar declines in inequality can be shown for other inequality statistics, such as the mean logarithmic deviation, the Theil Index, and the Atkinson family of inequality indices.

Finally, for many theoretical concepts of welfare (e.g. Atkinson's expected utility for the society, or Sen's real national income) it is possible to find an inequality index described above such that the welfare concept can be represented as GDP multiplied by one minus the inequality index. Since we can compute these inequality indices, we can show that because world inequality fell, welfare measured for the world as a whole grew even faster than world GDP did, and more than doubled over the period 1970-2006.

Figure 3. World gini inequality
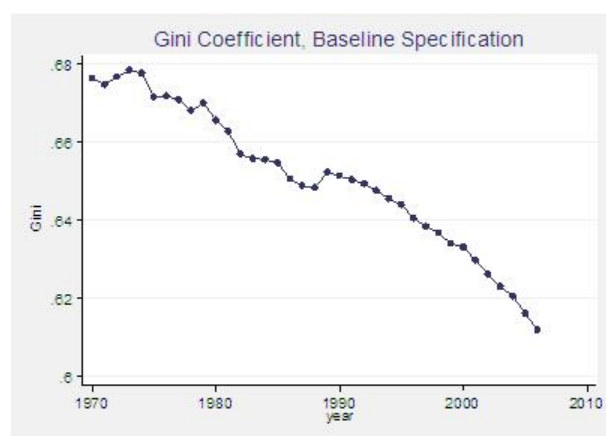
Another distinct advantage of using a parametric assumption is that we can correct one of the potentially largest problems of the survey data – the bias in the surveys arising from the underreporting of the very rich and the very poor.

The rich tend to not respond to surveys more systematically than middle-income people do (their houses are less accessible) and they may tend to underreport a larger fraction of their income as they have more incentives to hide.

The very poor, on the other hand, may be hard to interview (especially in poor countries) as they do not have identification or a fixed address. It may also be difficult to value the income of the very poor, as it often accrues in kind rather than in currency.

Underreporting may lead to understating poverty, if the poor or the rich are not counted, making the distribution seem more egalitarian than it is. Underreporting may instead lead to overstating it, if much

of the income of the poor is in kind, and not accounted for, then the poor seem to be poorer than they really are. Using the lognormality assumption, we can correct for misreporting by basing the distribution parameter not on all five quintiles, but on the middle three (specifically, on the ratios of each of the middle three quintile shares to their sum), since incomes in the middle quintiles should not be affected by problems in the tails of the distribution.

We find that underreporting is not very significant as the poverty series with and without the correction for underreporting are close together and underreporting actually overstates poverty. Hence, correcting for mismeasurement in the income survey data actually strengthens our conclusions.

Finally, our methodology allows us to conduct a thorough stress-test of our results to alternative assumptions. World poverty and inequality fall markedly for all the variations we try. In particular, we assess the impact of the World Bank's PPP revision that took effect in 2007. Under this revision, China's and India's GDP was reduced by 40% and 35% respectively compared with previous estimates. The revision has been criticised for estimating Chinese prices from urban data only which almost certainly overstates rural Chinese prices. Taking the PPP-revised estimates as given, however, we can rerun our procedure using them as our source of GDP. It turns out that while poverty rates increase under the new GDP (since China and India are 40% poorer), the amount of poverty reduction the world has experienced since 1970 actually *increases* – the world has been doing a better, not a worse job reducing poverty. This is because China is so rich now relative to the $1 a day poverty line that decreasing its GDP by 40% does not put so many more people into poverty, whereas in 1970, when China was among the poorest countries in the world, decreasing its GDP by 40% raised its poverty rate from about 74% to about 88%. Using the new World Bank numbers, we see that while the world looks like a somewhat worse place at any given point in time, its progress at reducing poverty has actually been even more impressive than we would have thought without the revision (Figure 4).

Figure 4. World Bank PPP revision
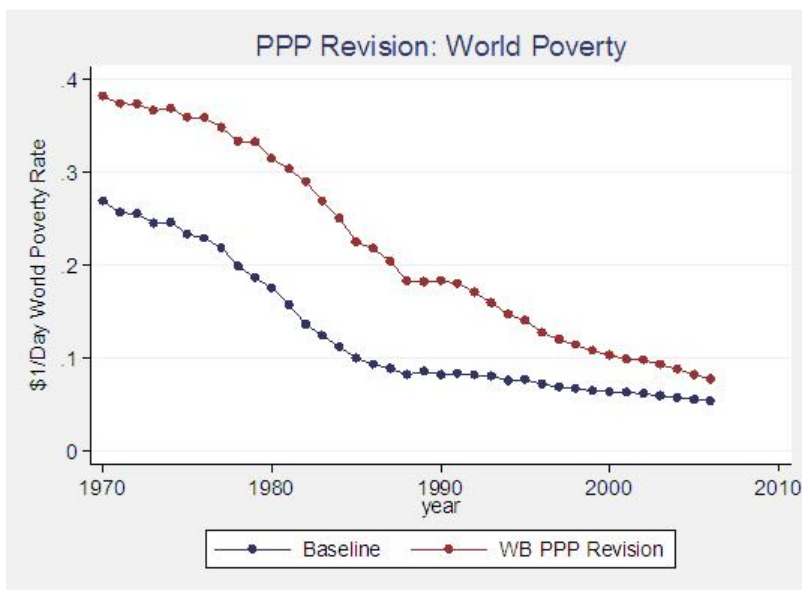
References

Gibrat, Robert (1931): Les inégalités économiques, Paris, Sirey.

Pinkovskiy, Maxim and Xavier Sala-i-Martin (2009), "Parametric Estimations of the World Distribution of Income", NBER Working Paper 15433.

Ravallion, Martin and Shaohua Chen (2008), "The World is Poorer than We Thought, but No Less Successful in the Fight Against Poverty", World Bank: Policy Research Working Paper 4703.

Sala-i-Martin, Xavier (2006), "The World Distribution of Income: Falling Poverty and...Convergence, Period," *Quarterly Journal of Economics*, vol. 121(2): 351-397.

UNU-WIDER (2008). UNU-WIDER World Income Inequality Database, Version 2.0c, May.

World Bank (2009). World Development Indicators, online database.

30

Topics: Poverty and income inequality
Tags: Income distribution, population, Poverty

VoxEU.org