

How Much Does Reducing Inequality Matter for Global Poverty?

Christoph Lakner

Daniel Gerszon Mahler

Mario Negre

Espen Beer Prydz



WORLD BANK GROUP

Development Data Group

Development Research Group

&

Poverty and Equity Global Practice

May 2019

Abstract

The goals of ending extreme poverty by 2030 and working toward a more equal distribution of income are prominent in international development and agreed upon in the United Nations' Sustainable Development Goals 1 and 10. Using data from 164 countries comprising 97 percent of the world's population, this paper simulates a set of scenarios for global poverty from 2018 to 2030 under different assumptions about growth and inequality. This allows for quantifying the interdependence of the poverty and inequality goals. The paper uses different assumptions about growth incidence curves to model changes in inequality and relies on the Model-based Recursive Partitioning machine-learning algorithm to model how growth in GDP is passed through to growth as observed in household surveys.

When holding within-country inequality unchanged and letting GDP per capita grow according to International Monetary Fund forecasts, the simulations suggest that the number of extreme poor (living below \$1.90/day) will remain above 550 million in 2030, resulting in a global extreme poverty rate of 6.5 percent. If the Gini index in each country decreases by 1 percent per year, the global poverty rate could reduce to around 5.4 percent in 2030, equivalent to 100 million fewer people living in extreme poverty. Reducing each country's Gini index by 1 percent per year has a larger impact on global poverty than increasing each country's annual growth 1 percentage point above the forecasts, suggesting an important role for inequality on the path to eliminating extreme poverty.

This paper is a product of the Development Data Group, the Development Research Group, and the Poverty and Equity Global Practice. It is part of a larger effort by the World Bank to provide open access to its research and make a contribution to development policy discussions around the world. Policy Research Working Papers are also posted on the Web at <http://www.worldbank.org/prwp>. The authors may be contacted at clakner@worldbank.org, dmahler@worldbank.org, mnegre@worldbank.org and eprydz@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

How Much Does Reducing Inequality Matter for Global Poverty?

Christoph Lakner
Daniel Gerszon Mahler
Mario Negre
Espen Beer Prydz*

JEL codes: I32, D31, O15

Keywords: Poverty, inequality, inclusive growth, simulation, machine learning

* All authors are with the World Bank. Negre is also affiliated with the German Development Institute. Contact information: clakner@worldbank.org, dmahler@worldbank.org, mnegre@worldbank.org, eprydz@worldbank.org. The authors wish to thank Shaohua Chen, Francisco Ferreira, La-Bhus Fah Jirasavetakul, Dean Joliffe, Aart Kraay, Peter Lanjouw, Christian Meyer, Prem Sangraula and Renos Vakis, as well as two anonymous referees for helpful comments and suggestions. The findings and interpretations in this paper do not necessarily reflect the views of the World Bank, its affiliated institutions, or its Executive Directors. Part of this work was funded by the UK Department for International Development through its Strategic Research Program (TF018888). This working paper is a substantially revised and updated version of Lakner et al. (2014). The earlier version focused on changes around the bottom 40% using a simple step-function growth incidence curve (GIC), whereas this paper considers more general distributional changes in the Gini index and more plausible functional forms of the GIC. Furthermore, this paper offers a more complete assessment of the potential tradeoffs between reducing inequality and increasing growth through the use of iso-poverty curves. Finally, this paper proposes a novel way to estimate the passthrough rate from GDP growth to growth in household survey income or consumption.

1 Introduction

Over the past two and a half decades, global extreme poverty has decreased rapidly. Since 1990, the share of the world population living below the extreme poverty line of \$1.90 per day has fallen from 35.6% in 1990 to 10.0% in 2015 (World Bank, 2018a). Against this backdrop, international development actors, bilateral development agencies and countries themselves have united around a goal of ‘ending’ extreme poverty by 2030. This goal has been defined as complete eradication (United Nations, 2014) or as reducing global extreme poverty to 3% of the world’s population (World Bank, 2014). Several bilateral development agencies, such as DFID and USAID, have also made such goals central to their focus and mission. At the same time, the development policy debate is increasingly paying attention to the level of inequality in countries around the world (International Monetary Fund, 2014; Ravallion, 2001; World Bank, 2016). As a result, the internationally agreed Sustainable Development Goals (SDGs) include both a goal to end poverty (SDG1) and a goal to reduce inequality within countries (SDG10).

We simulate global extreme poverty until 2030 under different scenarios about how inequality and growth evolve in each country. This serves to quantify the importance of reducing inequalities vis-à-vis increasing growth in achieving the goal of eradicating extreme poverty. Although previous papers have simulated poverty up to 2030, we offer four distinct contributions. *First*, we use micro data for 119 countries and grouped data for an additional 45 countries, allowing for an unprecedented data coverage of 97% of the world’s population. *Second*, we model the impact of distributional changes on future trajectories of global poverty by changing countries’ Gini index. The Gini index is arguably the most frequently used measure of inequality, and it makes for an intuitive way of modeling distributional changes which has direct policy relevance and conceptual simplicity. *Third*, since there are infinitely many ways in which a change in Gini indices can occur, we use different growth incidence curves to capture how inequality reductions may occur in an intuitive manner. *Fourth*, addressing the criticism that economic growth in national accounts is increasingly disconnected from income and consumption as observed in surveys (Ravallion, 2003; Deaton, 2005; Pinkovskiy & Sala-i-Martin, 2016), we utilize a novel machine-learning algorithm to estimate the share of economic growth passed through to income or consumption observed in surveys.

Our simulations suggest that the global poverty rate will remain around 6.5% in 2030 if growth is distribution-neutral and follows IMF forecasts. Under a scenario in which the Gini index of each country decreases by 1% per year, the global poverty rate falls to 5.4% -- equivalent to 100 million fewer people living in extreme poverty. Reducing each country’s Gini index by 1% per year has a larger impact on global poverty than increasing each country’s annual growth rate 1 percentage point (pp) above IMF forecasts. Even under the most optimistic scenarios we consider – where the Gini decreases 2% annually and the annual growth rate exceeds IMF forecasts by 2 pp – the poverty rate in Sub-Saharan Africa would remain around 20% in 2030 and the global target of 3% would not be met.

We simulate all changes in Gini indices at the national level, not globally. A pro-poor distributional change as simulated in this paper implies a fall in *within-country* inequality, but can be expected to have a more

mutated effect on *global* inequality, for which between-country differences matter greatly. One challenge with modeling the impact of changes in the Gini index on poverty is that there are infinitely many possible distributional changes resulting in the same change in the Gini index. If the change in the Gini index comes from redistributing resources from the wealthiest 1% to the middle class, poverty may remain unchanged in countries with moderate to low levels of poverty. If the change comes from instituting a basic income to all households, then a similar change in the Gini may completely eliminate poverty. Our baseline results are based on a linear growth incidence curve, but in a robustness check we use a convex growth incidence curve (GIC), which gives higher growth rates to the lowest percentiles compared to the linear version. With the convex functional form, a 1% annual decrease in the Gini in all countries has a larger impact on global poverty than a 2 pp higher annual growth in each country. In other words, the convex GIC further highlights the importance of reducing inequality for ending extreme poverty.

The literature has adopted several alternative approaches to model distributional changes in simulating global poverty trajectories. Some authors have simply imposed distribution-neutral growth, thus ignoring any future changes in within-country inequality (Birdsall et al., 2014; Karver et al., 2012; Hellebrandt and Mauro, 2015). Others have projected distribution-neutral growth but chosen initial distributions with different levels of inequality (Ravallion, 2013; Edward and Sumner, 2014). Other studies, which are most closely related to the approach taken by this paper, simulate additional distributional changes, by extrapolating the trend in the Q5/Q1 ratio (Edward and Sumner, 2014; Hillebrand, 2008; Higgins and Williamson, 2002), the Palma ratio (Chandy et al., 2013), or the income share of the bottom 40% (Ncube et al., 2014). A previous version of this paper used differences in growth rates of the bottom 40% and the mean to project poverty towards 2030 (Lakner et al. 2014), similar to Hoy and Samman (2015).

While our focus is on the impact of the distributional nature of future growth, we also develop our own baseline distribution-neutral growth scenarios. Two main approaches are used in the literature, which can produce quite different results for global poverty (Dhongde and Minoiu, 2013; Edward and Sumner, 2014). First, scenarios based on historical survey growth rates (e.g. Yoshida et al., 2014). Second, scenarios derived from national accounts either through growth models (Birdsall et al., 2014; Hillebrand, 2008), or projecting historical or forecasted growth rates into the future (Karver et al., 2012). Similar to our approach (explained in more detail in Section 4), Chandy et al. (2013) use Economist Intelligence Unit (EIU) and IMF's World Economic Outlook (WEO) growth rates adjusted to survey growth using factors from a cross-country regression. We base our projections on both country-specific historical growth rates and forecasted growth rates, adjusted for observed differences between household survey growth and national accounts growth. The distribution-neutral global poverty projections remain at around 6.5% in 2030 regardless of which growth scenario we use.

We model distributional changes and growth rates in GDP independently of each other. Although the famous Kuznets Hypothesis (Kuznets, 1955) would predict that higher growth in low-income countries would tend to increase inequality, the empirical support for this hypothesis is weak. Ferreira and Ravallion

(2009), for example, find no correlation between growth and changes in inequality in the developing world.²

The paper is structured as follows. Section 2 describes the conceptual framework for the simulations, while Section 3 describes the data and our method for implementing the simulations. Section 4 presents the results on global and regional poverty for different growth and inequality scenarios, while Section 5 presents robustness checks by using different growth incidence curves, poverty lines, and poverty measures. Section 6 concludes.

2 Conceptual Framework

In this paper, we model how changes in each country's Gini index impact poverty towards 2030. The Gini index is arguably the most frequently used measure of inequality, and it makes for an intuitive way of modeling distributional changes which has direct policy relevance and conceptual simplicity. One challenge with modeling the impact of changes in the Gini index on poverty is that there are infinitely many possible distributional changes resulting in the same change in the Gini index. To conceptualize this, we use GICs, and in particular restrict our focus to two functional forms of the GIC.³ Let y_i be the mean income of percentile group i (e.g. the bottom 1%) in the initial period. Final mean income y_i^* can be expressed as

$$y_i^* = y_i(1 + g_i) \tag{1}$$

where g_i is the growth rate associated with this percentile group. We define the GIC as the plot of g_i against the percentile group (p_i) in the initial period.

An intuitive and convenient way to allow the Gini to change is through a tax and transfer scheme introduced by Kakwani (1993) and further discussed by Ferreira and Leite (2003). This scheme involves an increase of everyone's income at a rate γ together with a tax and transfer scheme which taxes everyone at a rate τ and gives everyone an equal absolute transfer. As pointed out by Ferreira and Leite (2003), this is a type of Lorenz-convex transformation. They show that the transformed Lorenz curve is given by $L(p)^* = L(p) + \tau(p - L(p))$, where $L(p)$ is the original Lorenz curve, which is a function of the percentile p , where it is evaluated, and $L(p)^*$ is the post-transfer Lorenz curve. This transformation can be obtained by moving every point on the Lorenz curve upwards by an amount proportional to its vertical distance to the equidistribution (45-degree) line. The transformed Gini index can be readily obtained as $Gini(y)^* = (1 - \tau)Gini(y)$. In other words, the tax rate imposed, τ , is equivalent to the percentage change in Gini observed, α , such that $\tau = -\alpha$. This direct link between the tax-and-transfer scheme and the change in

² That said, the method we use to determine the share of growth in GDP that is passed through to the welfare vector observed in surveys is determined by an algorithm which takes inequality levels as one of its potential input variables. In other words, inequality is allowed to influence growth rates in welfare if there are empirical reasons for making such a connection. This turns out to be partially the case.

³ In Ravallion and Chen (2003), the GIC shows the growth rate of the income at a given percentile (e.g. the 10th percentile) between the initial and final period. In contrast, we compute the growth rate in the mean of a particular percentile *group*.

the Gini makes it a convenient way to model changes in the Gini. We can express the final incomes as a function of the initial income, mean income, and changes in the Gini:

$$y_i^* = (1 + \gamma)[(1 - \tau)y_i + \tau\mu], \quad (2)$$

where μ is the mean income in the initial period. Using (2) and (1), it can be shown that the corresponding GIC takes the following form:

$$g_i = (1 - \tau)(1 + \gamma) - 1 + [\tau(1 + \gamma)\mu] \frac{1}{y_i} \quad (3)$$

This GIC is a convex, decreasing function (when $\tau > 0$) along the percentile groups. It attributes high growth rates at lower percentiles, while it becomes flatter at higher percentiles. It is decreasing throughout, meaning that the growth rate will be lowest for the richest percentile groups.

Another way of simulating a change in the Gini index uses a linear GIC. Such a GIC takes the following form:

$$g_i = \delta - \theta p_i \quad (4)$$

Substituting (4) into (1), we can obtain the following expression for the income of percentile group i in the final period

$$y_i^* = (1 + \delta)y_i - \theta y_i p_i \quad (5)$$

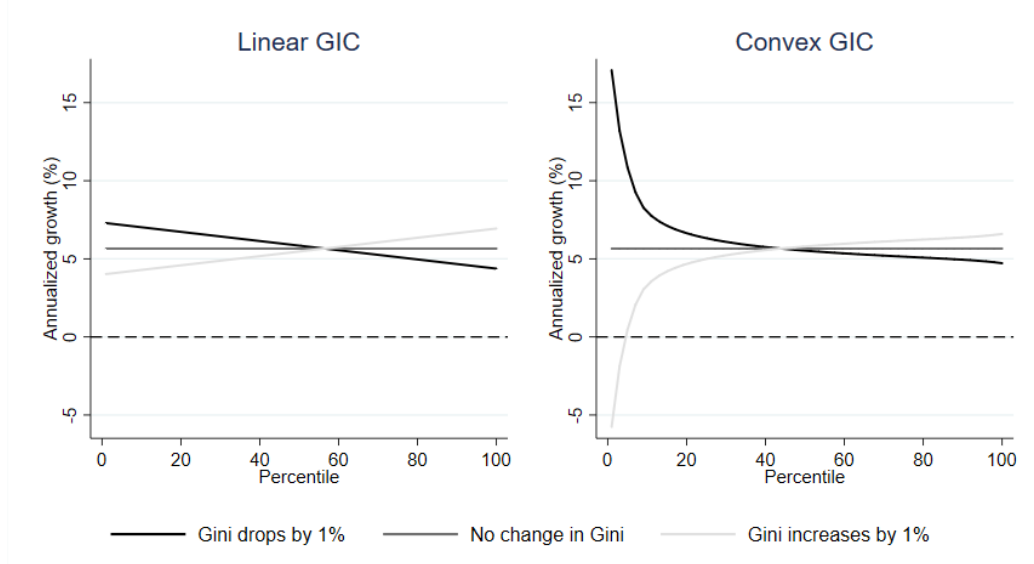
This linear GIC can be obtained by taxing everyone in proportion to both their income and rank – the poorest person is taxed at a rate of θ and the tax increases proportionally with the rank – combined with a transfer where every person receives a share δ of their income. Unlike the convex GIC, whose central parameter is directly related to percentage changes in the Gini index, there is no functional relationship between the percentage change in the Gini index, α , and the parameters of the linear GIC. We thus use an algorithm that iteratively changes the slope of the GIC until it matches the desired α .

To illustrate how the convex and linear GICs could look in practice, we use the welfare distribution from a survey in Côte d'Ivoire from 2015. From 2015 to 2016 IMF data suggest that GDP per capita in Côte d'Ivoire grew by 5.6%. Figure 1 explores how this growth can be distributed if inequality stays unchanged or if the Gini increases or decreases by 1% (ignoring for the moment that only part of this growth is passed through to the consumption observed in surveys). The initial Gini in Côte d'Ivoire was 41.5, meaning that a 1% drop ($\alpha = -0.01$) would bring the Gini to 41.1, while a 1% increase ($\alpha = 0.01$) would bring it to 41.9.

Lowering the Gini by 1% does not have to impose a large cost on the top of the distribution. Because of the large income share of the top of the distribution, the reduction in the growth rate of the wealthiest individuals necessary to ensure that the bottom grows substantially faster than the mean is relatively small. For example, in the case of Côte d'Ivoire, a convex growth incidence such that the Gini decreases by 1% means that households at the 10th percentile grow 2.5 pp faster than the mean, yet only reduces the growth at the 90th percentile by 0.5 pp.

In our baseline simulations we use linear GICs. There are three reasons for this choice: First, it is probably the simplest realistic pro-poor GIC that can be constructed. Second, it constitutes a relatively conservative pro-poor distributional change, in contrast to the convex GIC, which may provide a too optimistic picture of how reducing inequality affects poverty. Finally, in contrast to the convex GIC, it can easily be implemented for increasing Gini indices as well. A challenge with using convex GICs is that certain large increases in the Gini can only be implemented if the poorest households attain a negative income level. In those cases, the best solution may be to constrain the income levels to be zero, implying that the Gini does not increase as much as desired.

Figure 1: Different growth incidence curves compatible with same change in the Gini index



Note: Growth incidence curves (GICs) drawn using data from Cote d'Ivoire 2015 under different assumptions about how much inequality changes, and in what manner inequality changes. The mean is assumed to grow at 5.6%, according to data from the WEO.

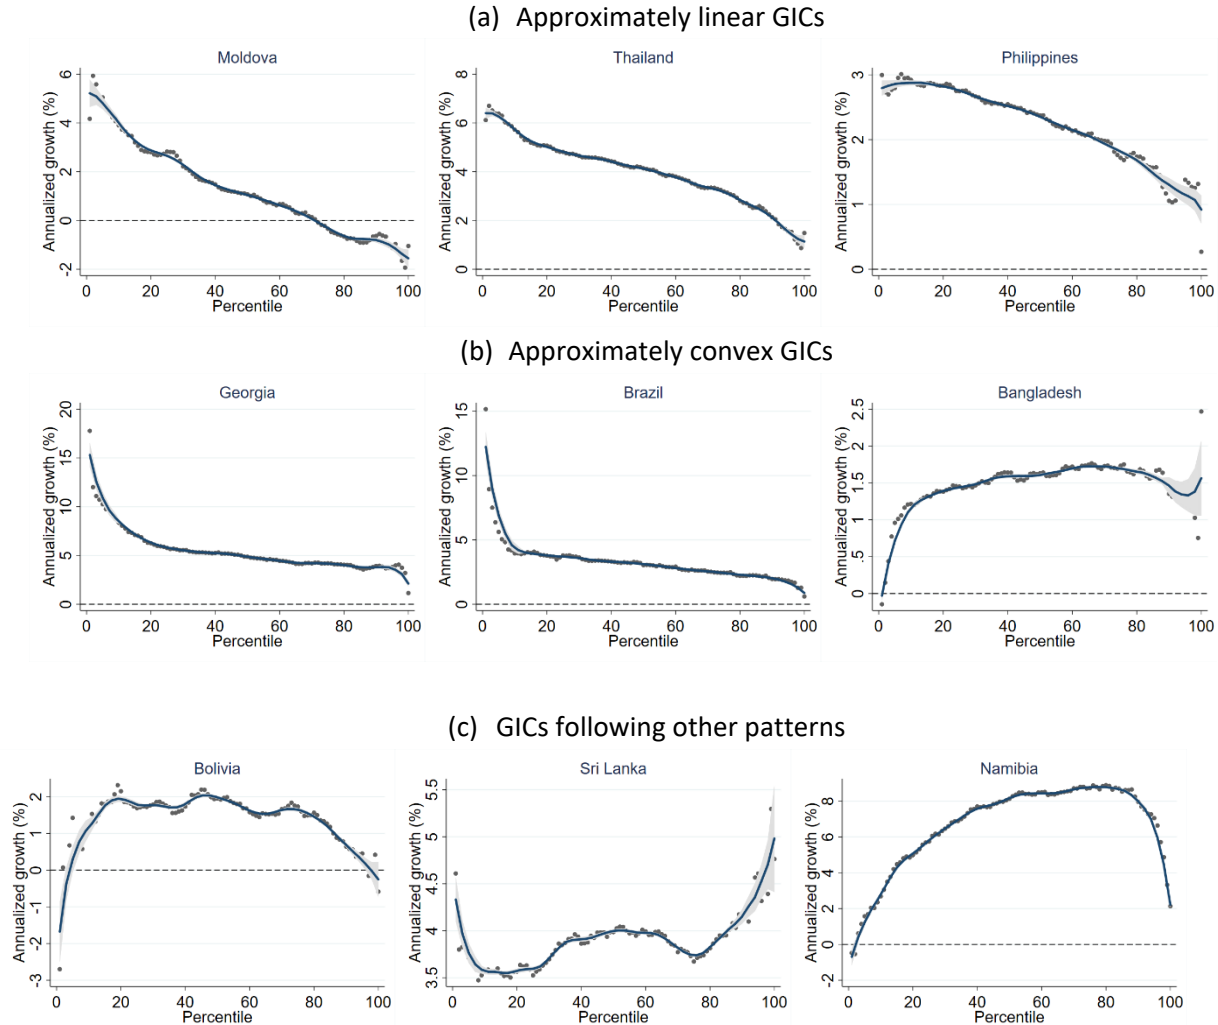
Nonetheless, the convex GIC has some advantages: First, it intuitively relates to public policy, as it represents what would happen to poverty if a linear tax of rate τ were implemented with a lump sum transfer. Second, it is analytically related to changes in the Gini index, allowing for a direct link with the measure of distributional change we are looking at. Third, it is directly linked to differences in growth rates of the bottom 40% and the entire distribution, also called shared prosperity, which is the first target of the SDG on inequality.⁴ For these reasons, we will use convex GICs as a robustness check.

A worthwhile question to ask is whether these GICs are observed empirically. Using the World Bank's Global Shared Prosperity Database (World Bank, 2018b), which provides a list of 259 spells with a

⁴ The percentage change in the Gini index and the shared prosperity premium (the difference in growth of the bottom 40% and the mean, denoted m) are related as follows: $\alpha = \frac{m}{(1+\gamma)(\frac{0.4}{s_{40}}-1)}$, where s_{40} is the income share of the bottom 40%. Hence, for a given income share of the bottom 40% and overall growth rate, there is a linear connection between the size of the tax rate, the percentage change in the Gini index, and the shared prosperity premium. For more details on the formal relationship between the convex growth incidence curve and shared prosperity, see the appendix of the earlier version of this paper (Lakner et al. 2014).

comparable welfare aggregate and surveys that lie about 5 years apart, we can explore how GICs for these countries look in practice. Figure 2 shows examples of GICs that look approximately linear, GICs that look approximately convex, and GICs that follow different shapes. Based on these patterns, we believe there are sufficient empirical examples of the two types of GICs that we will focus on in this paper to make them relevant.

Figure 2: Empirically Observed Growth Incidence Curves



Note: Empirically observed growth incidence curves using the surveys in World Bank (2018b).

A good alternative to using a theoretically defined GIC would be to impose one that has been observed in practice, e.g. for the same country, a best performer in the region, etc., as done in World Bank (2015). Yet, this does not provide a sense of the magnitude of the distributional change required, which our paper attempts to specify. It is also problematic for the many countries that lack comparable data over time, preventing historical GICs to be created.

3 Data and Methodology

3.1 PovcalNet

To predict poverty in 2030, we rely on the surveys used in PovcalNet, which contains the World Bank's official country-level, regional and global estimates of poverty. Most of the data in PovcalNet come from the Global Monitoring Database (GMD), which is the World Bank's repository of multitopic income and expenditure household surveys used to monitor global poverty. PovcalNet contains more than 1,500 surveys from 164 countries covering 97% of the world's population. The data available in PovcalNet are standardized as far as possible but differences exist with regards to the method of data collection, and whether the welfare aggregate is based on income or consumption. By relying on the PovcalNet database, we ensure consistency with the official numbers used by the World Bank and United Nations for monitoring poverty, inequality and related goals.

For 119 of the countries, housing 64% of the world's population, micro data are available. For an additional 35 countries, or 14% of the world's population, mostly comprising the high-income world, grouped data of 400 bins are available. For the purposes of these projections, we treat the bins as microdata. Finally, for China and nine other countries constituting about 20% of the world's population, only decile or ventile shares and the overall mean are available. This concerns Algeria, China, Guyana, St. Lucia, Macedonia, Suriname, Turkmenistan, Trinidad & Tobago, Venezuela and Zimbabwe. For these countries, based on the decile shares and means reported in PovcalNet, we use a lognormal Lorenz curve to generate a distribution of 10,000 points for each country. The use of a parametric Lorenz curve is very similar to what is done in PovcalNet to calculate poverty when micro data are not directly available.⁵

3.2 Growth scenarios

Our starting point is the latest available survey in each country or, if the latest available survey was conducted prior to 2015, the welfare distribution the World Bank used to measure poverty for the country in 2015, which is the latest year with global poverty estimates at the time of writing.⁶ The median year of the latest survey is 2015, but the range spans from 1992 to 2017. Before implementing various inequality scenarios, we bring the welfare aggregates from these surveys to 2018. To do so, we multiply each household's welfare with (one plus) a fraction of the observed or forecasted growth rate in real GDP per capita. That is, we assume that the growth that occurred between 2015 (or the year of the latest survey) and 2018 in each country was distribution-neutral. We assume that only a fraction of the growth is passed through to the welfare vector, following the method described in the next subsection.⁷

⁵ From two parametric Lorenz curves – the General Quadratic and the Beta Lorenz – PovcalNet chooses the one with the best fit. Shorrocks and Wan (2008) suggest that a lognormal functional form fits better. Minoiu and Reddy (2014) show that for global poverty estimates a parametric Lorenz curve should be preferred to estimating kernel densities. We use the ungroup command included in the DASP Stata Package (Abdelkrim and Duclos, 2007) to fit a lognormal Lorenz curve. This command implements the Shorrocks and Wan (2008) approach which ensures that the fitted Lorenz curve matches the observed shares.

⁶ The method that PovcalNet uses to bring up the surveys to a common reference year is described in Appendix A of World Bank (2018a).

⁷ We use GDP throughout for consistency across countries, while PovcalNet chooses between growth in terms of GDP and household final consumption expenditure in national accounts.

Our preferred source of growth data is annualized growth in real GDP per capita from national accounts, as reported in the World Development Indicators (WDI). When such data are not available for the whole period, we complement it with annualized growth in real GDP per capita from the WEO. In two economies, data are also missing from WEO (Syria, and the West Bank and Gaza). In these cases, we rely on GDP data from the Economist Intelligence Unit.

Beyond 2018, we use three different growth scenarios: (1) that each country grows according to its annualized growth rate from national accounts for the last 10 years for which we have data (2007-2017); (2) that each country grows according to its annualized growth rate from national accounts for the last 20 years (1997-2017); and (3) that each country grows according to its annualized projected growth rate from 2018-2023, which currently is the last year for which WEO has growth projections. The simulations relying on the 20-year historic growth rates (1997-2017) may be optimistic, as the rapid growth experienced in the early 2000s is showing signs of slowing down. For example, Rodrik (2014) suggests that the rapid growth experienced by emerging economies in recent decades is unlikely to persist indefinitely and that convergence will slow down in coming decades.

3.3 The relationship between GDP/capita growth and welfare growth from surveys

A challenge with using growth rates of GDP/capita is that prior evidence has shown that only a fraction of growth observed in national accounts is passed through to the growth observed in household surveys (Ravallion, 2003; Deaton, 2005; Pinkovskiy & Sala-i-Martin, 2016). Estimating this fraction across our entire sample is fairly straightforward. One would simply regress the annualized growth in the survey means on the annualized growth in real GDP/capita, under the constraint that the intercept is zero, $g_{survey} = \beta * g_{GDP/capita} + \varepsilon$, and use β as the fraction of growth in GDP/capita that is passed-through to the welfare observed in surveys. Using a sample of 1,351 spells for which we can calculate changes in the survey mean suggests that $\beta = 0.83$. Each spell relies on two adjacent surveys from the same country with welfare measured in the same way (either income or consumption).

Yet there is no reason to believe that β is constant across different contexts. It may differ by geographical region, by income level, by whether income or consumption is used as the input into poverty measurement, over time, etc. Although interactions for these additional covariates can easily be accommodated in the equation, it is not clear which variables should be used to define the interactions and using all possible interactions will likely overfit the data. Applying a selected number of interactions is common practice in adjusting between household survey growth and national accounts growth rates (see for example Birdsall et al., 2014; Chen and Ravallion, 2010; and Chandy et al., 2013), but it is not entirely clear on what basis the interactions were selected.

To circumvent this issue, we apply a machine learning algorithm, Model-based Recursive Partitioning, to determine when there is reason to believe that the passthrough rate varies in different contexts (Zeileis et al. 2008). This algorithm can take as input all potential variables that might matter for the passthrough rate. In our case, as input variables we use geographical region (we use two versions, the official World Bank geographical regions, and the regions from PovcalNet, where most high-income countries form a separate region), a dummy for whether consumption or income is used, mean consumption, median consumption, the Gini index, population, GDP/capita, and the year of the survey.

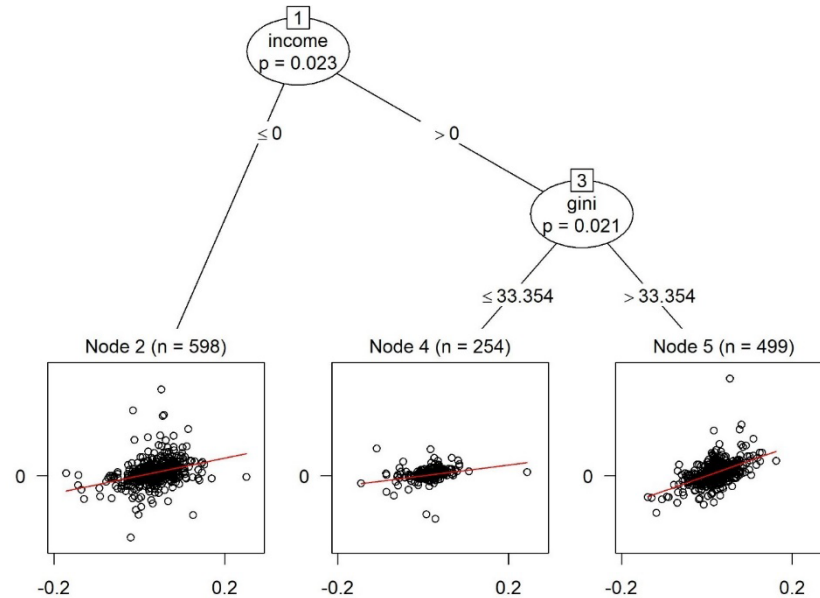
In short, the algorithm tests if the passthrough rate (β) depends on a particular input variable. When it does, it splits the sample into two based on this variable, and estimates the passthrough rate for each subsample separately. This continues iteratively in the two subsamples, such that the algorithm tests for evidence in favor of different passthrough rates in each of the subsamples, and splits the sample into two if passthrough rates differ sufficiently. The algorithm is a variant of classification and regression trees, pioneered by Breiman et al. (1984).

In more detail, the algorithm works in the following manner:

1. Run the regression $g_{survey} = \beta * g_{GDP/capita} + \varepsilon$ on all relevant data.
2. Add interactions between $g_{GDP/capita}$ and each of the input variables separately, and conduct Wald tests indicating whether the interaction coefficient(s) are statistically significant.
3. If the lowest p-value of these interaction coefficients (after adjusting for multiple hypothesis testing) is less than 0.05, then the variable with the lowest p-value is chosen as a splitting variable. If the lowest p-value is greater than 0.05, no split is made, and the algorithm stops (suggesting there is no evidence in favor of passthrough rates differing by context).
4. Split the sample into two using the splitting variable. If the splitting variable is not binary, meaning there is more than one way of splitting the sample into two, all possible splits are tried out (respecting monotonicity for continuous and ordered variables), and the split that results in the greatest rejection of equality of the passthrough rates is chosen, and the sample is split into two.
5. The algorithm is repeated from the beginning by applying it to observations in each of the two subsamples separately.

Figure 3 shows the results of this procedure using our data at hand. There is significant evidence in favor of the data type mattering for passthrough rates (0=consumption, 1=income). Observations using income have a passthrough rate of 0.99, while observations using consumption have a passthrough rate of 0.73. With a p-value of 0.023, we can reject that the coefficient is identical for the two subgroups at a 5% level. For observations using consumption, there is no variable which significantly yields different passthrough rates. For the observations using incomes, the Gini index matters for determining the passthrough rate. Cases with a Gini above 33 have a passthrough rate of 1.22 while observations with a Gini below 33 have a passthrough rate of 0.44. Table 1 contains more details on the Wald tests, the splits conducted, and the associated passthrough rates.

Figure 3: Decision tree of passthrough rates



Note: Results of using model-based recursive partitioning to determine when passthrough rates differ in various contexts. The figure should be read from the top down. The circles show the variable for which passthrough rates differ significantly and the p-value associated with the Wald test. The square boxes show the resulting regression plot and the fitted line. The income variable takes the value 0 if consumption is used and 1 if income is used.

Table 1: Details on decision tree algorithm

Node	Obs.	β	p-values from Wald tests									
			Inc/co ns	Gini	Media n	Mean	GDP	World Bank region	Povcal Net region	Year	Popul ation	Headc ount
1	1351	0.83	0.02	0.22	0.90	0.16	0.11	0.99	0.99	0.99	1.00	1.00
2	598	0.71	----	1.00	0.17	0.15	0.19	1.00	1.00	0.70	1.00	0.74
3	753	0.99	----	0.02	0.64	0.95	0.94	1.00	0.91	0.15	0.93	0.40
4	254	0.44	----	1.00	1.00	1.00	1.00	1.00	1.00	0.98	0.99	1.00
5	499	1.22	----	0.98	0.99	1.00	1.00	1.00	0.75	0.46	0.86	0.93


Note: The table shows the number of observations in each node ((sub)sample) of the tree, and the passthrough rate for observations in each node. The columns to the right show the p-values (adjusted for multiple hypothesis testing) from the tests exploring if passthrough rates vary by the variable in question in each particular node. Elements in bold show the p-values that govern the splits in the tree. The different nodes are defined as follows: (1) Full sample; (2) only consumption surveys; (3) only income surveys; (4) only income surveys with Gini ≤ 33.35 ; (5) only income surveys with Gini > 33.35 . "----" indicates that no test can be conducted since there is no variation in the input variable in question in the particular subsample.

3.4 Inequality scenarios

With the growth rates and passthrough rates in place, we consider five different scenarios for changes in the Gini index; that it changes by -2%, -1%, 0%, 1% and 2% per year. If a country starts with a Gini index of 0.40 in 2018 (which is close to the median Gini of the latest survey for each country), under our five different scenarios, it would end up with a Gini of 0.31, 0.35, 0.40, 0.45 and 0.51, respectively. Evaluating the plausibility of these Gini changes is difficult due to the lack of comparable data across countries over time. If the methodology for computing the consumption aggregate changes from one survey to the next, then the observed Gini might change drastically despite the true Gini remaining unchanged.

Utilizing the World Bank’s Global Shared Prosperity Database (World Bank, 2018b), we can look at 259 spells with a comparable welfare aggregate and surveys that lie about 5 years apart. The histogram of annualized percentage changes in the Gini index from these 259 spells is plotted in Figure 4. The histogram suggests that annualized changes in the range of -2% to 2% are not unlikely. In nearly half of the spells, the annual change in the Gini in absolute terms is 1% or greater, and in 15% of the cases the annual change in the Gini in absolute terms is 2% or greater. While the latter clearly are optimistic and pessimistic scenarios, especially if they are assumed to persist for 12 years, they are not unprecedented in past spells.

Figure 4: Observed annualized changes in the Gini compared with changes used in simulations



Note: Histogram of empirically observed annualized changes in the Gini using the spells from the surveys in World Bank (2018b).

3.5 Estimating global poverty

Armed with growth rates, passthrough rates, and changes in the Gini index, using the linear or convex growth incidence curve, we can project the welfare distribution in each country towards 2030. To project the distribution, we use the povsim simulation tool (Lakner et al. 2014).⁸

In order to derive global poverty rates, a few more pieces are needed. First, we need consumer price indices (CPI) and purchasing power parity (PPP) exchange rates to convert the national welfare aggregates into constant USD that have been adjusted for international price differences. To that end, we rely on the

⁸ Povsim can be installed in Stata by typing “net install povsim, replace from(<http://eprydz.com/povsim/>)”.

data used by PovcalNet. Most CPIs are from the IMF's International Finance Statistics, while most PPP exchange rates are from the International Comparison Program (PPPs for household final consumption expenditure). More details on the price data used are available in Lakner et al. (2018) and Atamanov et al. (2018). Second, we need population data to aggregate poverty estimates across regions and globally. We use annual population projections for each country from the World Bank.⁹ Finally, to arrive at regional and global poverty rates, we also need estimates for the 3% of the world for which we have no distributional data. In these cases, we follow the aggregation method used by Chen and Ravallion (2010) and deployed by PovcalNet, which assumes regional poverty rates for countries without a poverty estimate.

4 Results

This section presents the results from the simulations described above. First, we show distribution-neutral poverty projections towards 2030, both at the global and regional level, focusing on the poverty rate measured at \$1.90/day in 2011 PPPs.¹⁰ Second, we explore what would happen if growth or inequality changes in a positive or negative direction, and the trade-off between increased growth and reduced inequality.

4.1 Impacts on Poverty: Global and Regional Trajectories to 2030

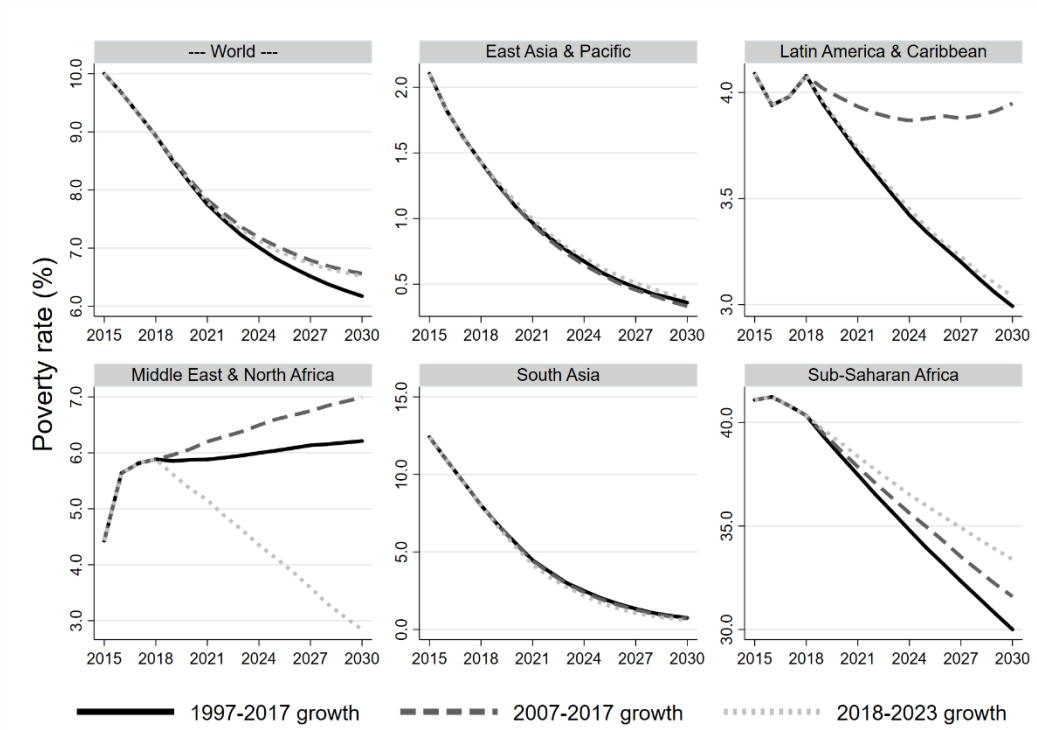
Figure 5 presents our simulated trajectories for the global poverty rate to 2030 for the three different distribution-neutral growth scenarios: that countries follow their growth patterns of the past 10 years, of the past 20 years, or that they follow the growth projections from the WEO.

All scenarios put the global poverty rate in 2030 in the range of 6-7%. The scenario using historical growth rates from 1997-2017 is a bit more optimistic, due to the high growth rates at the turn of the century. While the different growth scenarios do not matter much at a global level, there are starker differences at the regional level. In Latin America & the Caribbean, using growth rates from the past 10 years results in almost no decrease in poverty – a pattern mostly attributable to poor recent growth in Venezuela -- while the other two scenarios decrease poverty substantially towards 2030. In the Middle East & North Africa, both historical growth scenarios yield increasing poverty rates towards 2030 while using the WEO growth projections suggests that poverty will be halved by 2030. The global poverty rate is largely driven by Sub-Saharan Africa, which in all three scenarios has poverty rates at or above 30% in 2030, while the other regions of the world have rates below 7% (the vertical axis differs across regions in Figure 5).

⁹ The projections and estimates are available at <https://datacatalog.worldbank.org/dataset/population-estimates-and-projections>.

¹⁰ See Ferreira et al. (2016) for a description of how the \$1.90 international poverty line has been defined.

Figure 5: Global and regional distribution-neutral poverty projections to 2030

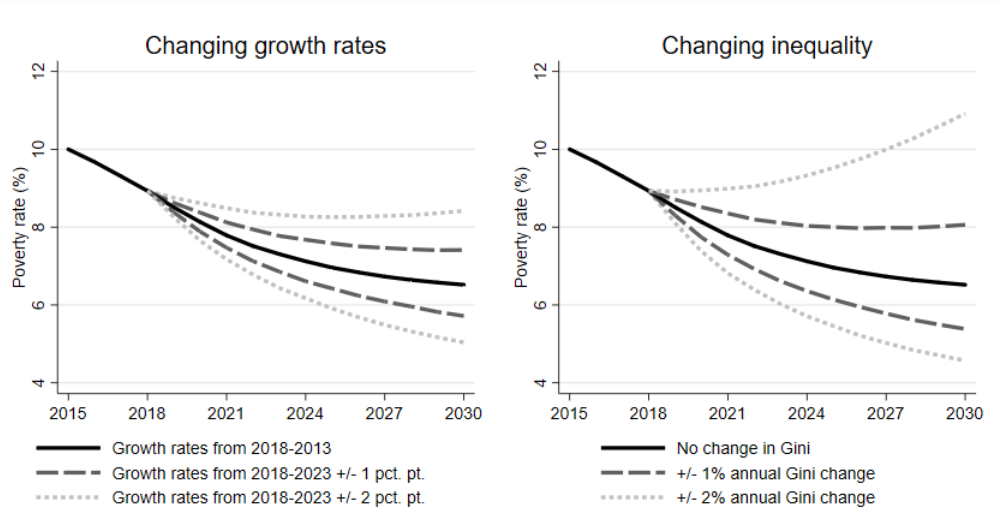


Note: Projected global and regional poverty rates measured at \$1.90/day in 2011 PPPs assuming distribution-neutrality under three different growth scenarios: countries follow their growth patterns of (1) the past 10 years, (2) the past 20 years, or (3) the growth projections from the WEO.

Next, we look at how changing the income distribution or the growth rates impacts global poverty. We focus on the 2018-2023 growth rate scenarios (i.e. the WEO forecasts), and simulate the change in poverty if each country’s annual growth rate is 1 or 2 pp higher than this growth rate. In addition, we consider simulations if each country’s Gini index decreases or increases by 1% or 2% per year using linear growth incidence curves. Results are shown in Figure 6.

Decreasing the Gini index by 1% annually in each country has a larger impact on poverty than increasing growth 1 pp above forecasts, and in general the projections are quite sensitive to changes in the Gini index. Under the same growth scenario, the global poverty rate could be between 4% and 11%, depending on the distributional nature of that growth.

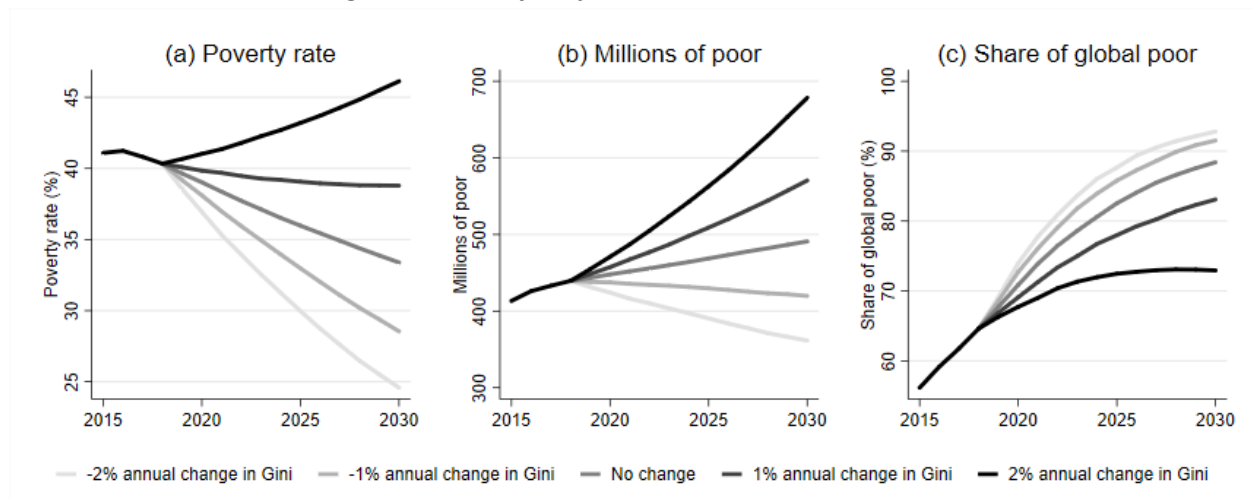
Figure 6: Simulations of global poverty under different growth and Gini scenarios



Note: Projected global poverty rate measured at \$1.90/day in 2011 PPPs assuming that countries (1) exceed or fall behind the growth projections from the WEO by 1 or 2 pp annually (left panel), or (2) follow the WEO projections exactly but reduce/increase their Gini index by 1 or 2% annually (right panel).

Changes in the Gini index are particularly relevant for Sub-Saharan Africa, where the poverty rate fluctuates from 24% to 46% depending on the distributional scenario (Figure 7a). Due to rapid expected population growth, only the scenarios that lower inequality are expected to decrease the number of poor in Sub-Saharan Africa (Figure 7b). Since the inequality-reducing scenarios rapidly reduce poverty in other regions, the share of the global poor that live in Sub-Saharan Africa actually increases under these scenarios. More than 90% of the global poor would reside in Sub-Saharan Africa by 2030 if all countries experience a fall in inequality.

Figure 7: Poverty Projections in Sub-Saharan Africa



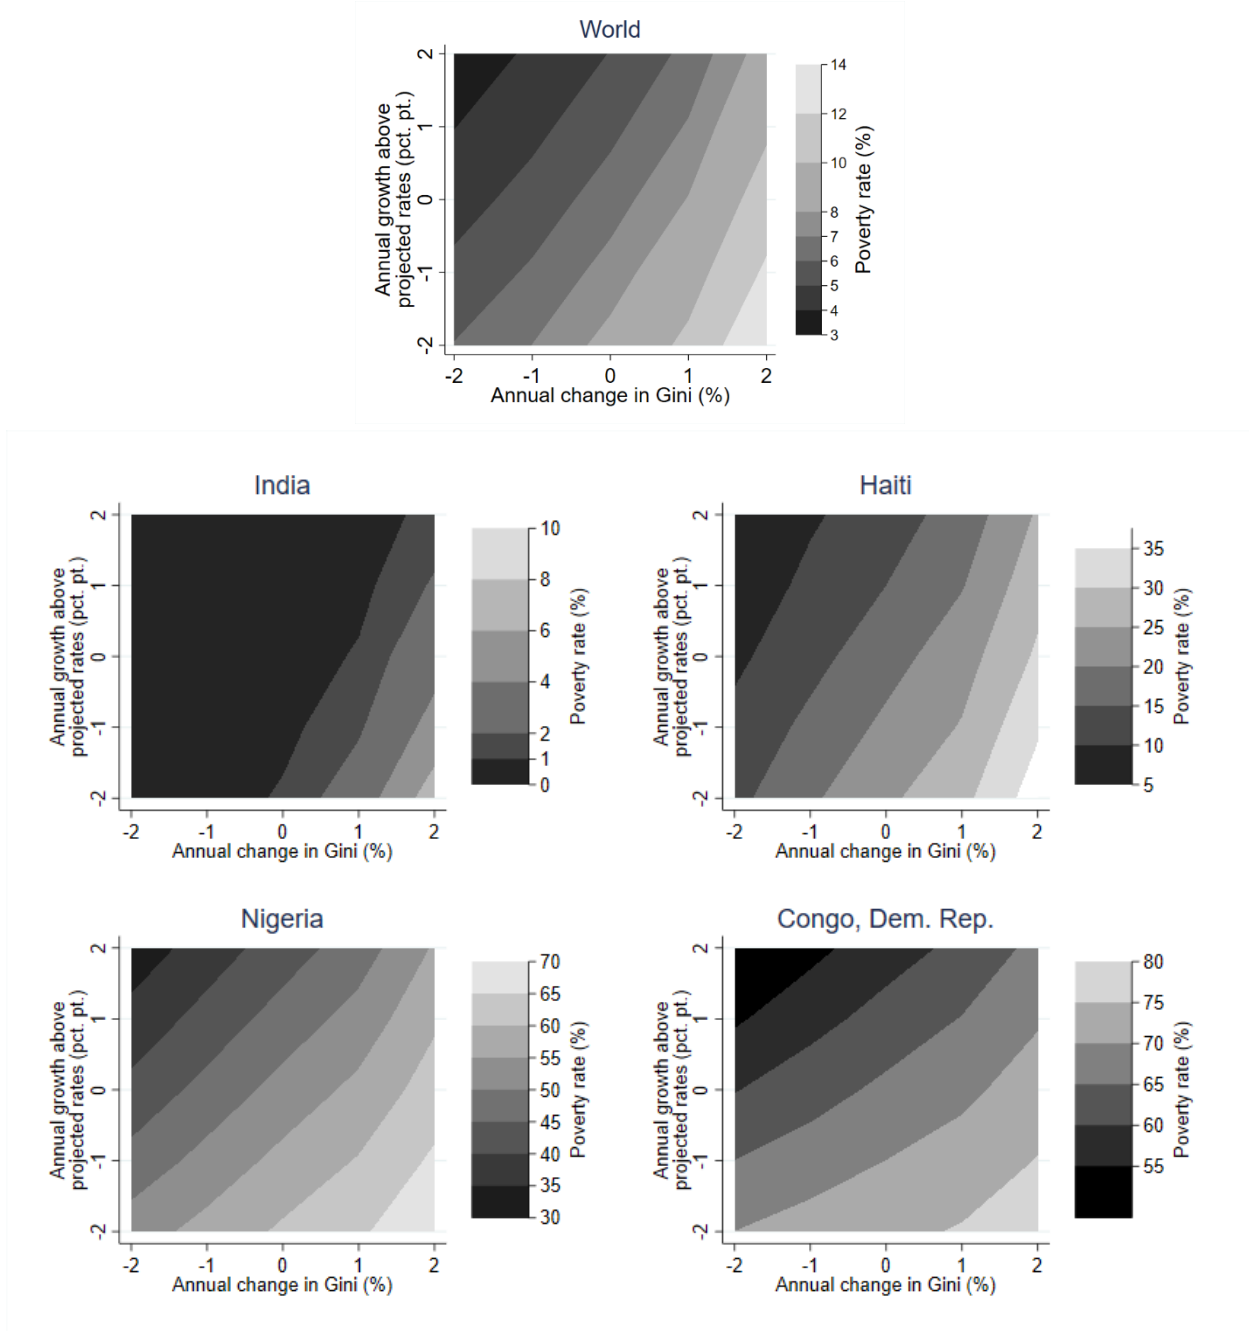
Note: Projected poverty rates in Sub-Saharan Africa measured at \$1.90/day in 2011 PPPs assuming that countries follow the growth projections from the WEO under five different scenarios about how inequality will change in each country.

The combinations of scenarios changing the Gini index and making the growth rate higher or lower than WEO projections allows for the creation of iso-poverty curves. These curves, introduced by Ferreira and Leite (2003), show combinations of inequality changes and growth changes resulting in the same level of poverty, as shown in Figure 8. The flatness of the curves illustrates the relative role of growth and inequality in shaping poverty rates.

At a global level, the curves are somewhat vertical, suggesting that reducing the Gini index with a linear GIC is more impactful than exceeding growth forecasts. This pattern varies greatly by country. For countries with low poverty rates, the picture is mostly the same, and changing the Gini generally has a greater effect than exceeding growth forecasts. For countries with high poverty rates, the opposite occurs. In these cases, where the initial poverty rate may be above 50%, inequality-reducing growth might even increase the poverty rate, as the ones on the margin of being poor will have resources transferred to the very bottom of the distribution. In the Democratic Republic of Congo, for example, the poverty rate is less responsive to changing the Gini index than to exceeding growth forecasts.

These conclusions are tied to the set-up we have explored. If we used higher poverty lines, other countries would present a similar pattern to that seen in the Democratic Republic of Congo. Conversely, if we use measures of poverty that account for the depth and severity of poverty, improving the conditions of the bottom of the distribution may unambiguously be beneficial. Finally, the global pattern of inequality being more important than growth is influenced by our choice of growth incidence curve. In the next section we will explore the robustness of the results to the choice of alternative poverty lines, poverty measures, and GICs.

Figure 8: Global and country-specific iso-poverty curves, 2030




Note: The figure shows different combinations of changes in the Gini index and exceeding/falling behind growth forecasts that results in the same poverty rate globally and for four selected countries. The flatter the curves, the more growth matters relative to reducing inequality.

5 Robustness Checks

Our results thus far have used a linear GIC. This placed a limit on the simulated growth rates for the poorest individuals. If a convex GIC is used instead, for countries with low and moderate poverty rates, the bottom of the distribution experiences large shifts in their welfare. To check the sensitivity of our results to our choice of GIC, we implement the changes using a convex GIC as well. The resulting global ISO-poverty curve is shown in panel (b) of Figure 9. Compared to our original iso-poverty curve, reproduced in panel (a), using a convex GIC increases the impact of Gini changes on poverty reduction, as shown by the iso-poverty curves becoming steeper. Now a 1% annual reduction in the Gini matters as much as exceeding growth forecasts by 2 pp annually, as both bring the global poverty rate in 2030 to about 5%.

Figure 9: Global iso-poverty curves, 2030 under different assumptions



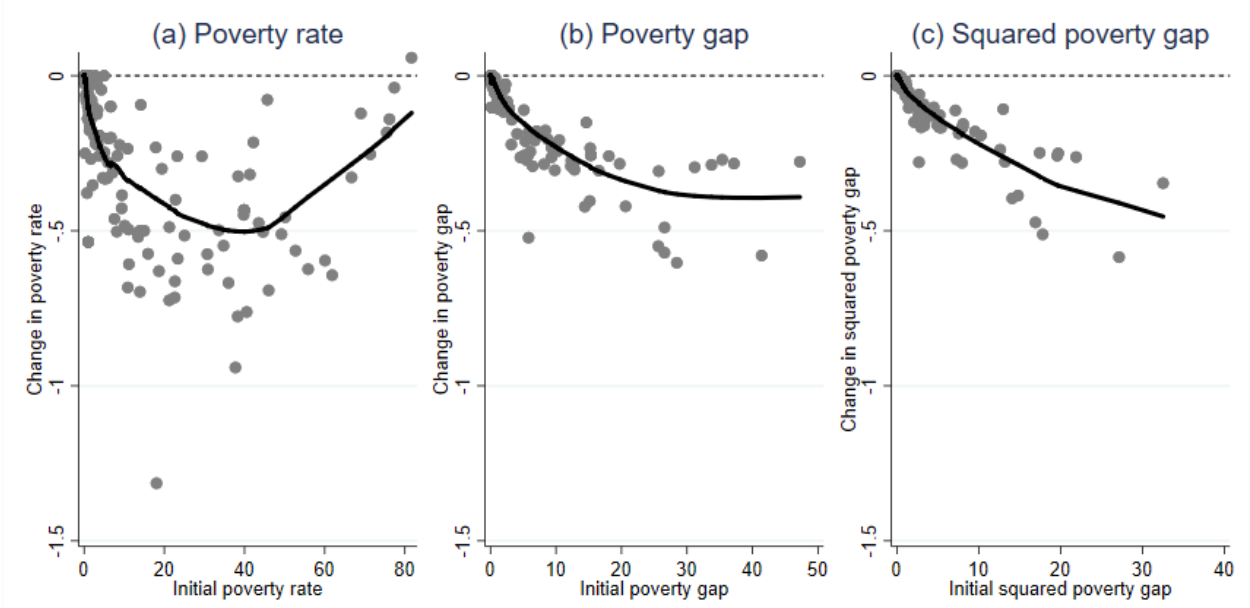
Note: The figure shows the global iso-poverty curve in 2030 under our baseline assumptions (panel a), and under five different robustness checks. Panel b uses convex GICs rather than linear GICs (all other panels use linear GICs). Panel c and d use higher poverty lines, than the \$1.90. Panel e and f use different poverty measures, the poverty gap and the squared poverty gap, respectively.

Next, we use higher poverty lines. Specifically, we use the poverty lines of \$3.2 and \$5.5, which are official higher poverty lines of the World Bank. These lines are constructed to reflect typical national poverty lines in lower- and upper-middle income countries, respectively (Jolliffe & Prydz, 2016). Under both lines, a 1% annual decline in the Gini still has a larger impact on global poverty than exceeding growth forecasts by 1 pp per year. This is despite of the fact that almost half of the world lived below \$5.50 in 2015 (World Bank, 2018a).

Finally, we use different poverty measures. The headcount ratio, which all our results thus far were based upon, is insensitive to the distributional differences among the poor, i.e. it does not value how far below the poverty line the poor fall. This may be an important measure to consider in countries with high poverty rates, where an inequality-reducing simulation may transfer resources from the marginally poor to the very poor. When using poverty measures that account for the depth and severity of poverty, FGT1 and FGT2 (Foster et al., 1984), the GIC becomes steeper, meaning that changes in the Gini have an even larger impact on poverty reduction, relative to higher growth (panel e and f).

The impact of reducing the Gini index varies with the initial poverty rate, as well as the shape of the GIC, the income distribution, and the measure of poverty. The impact of reducing the Gini by 1% is plotted against initial poverty levels in Figure 10 for linear GICs using three different measures of poverty. The figures are drawn for the change in predicted poverty from 2018 to 2019 assuming zero growth to abstract from differences in growth rates across countries.

Figure 10: Impact of reducing Gini by 1%, by poverty measure and poverty level



Note: Figures show the one-year change in poverty measures assuming the Gini decreases by 1%, zero per capita growth and a linear GIC.

The initial level of poverty matters for the impact of a fall in the Gini index on the poverty rate. The relationship takes a U-shape where the reduction in the poverty rate at first increases with the initial poverty rate, attains its maximum impact with poverty rates of about 40%, and then decreases (panel a).

For very high poverty rates, reducing the Gini increases poverty. Hence, there may be a certain trade-off between decreasing the poverty rate and decreasing inequality for very poor countries. Panels b and c show that reducing inequality almost unambiguously decreases both the poverty gap and the squared poverty gap even for high initial headcount ratios. This indicates that the trade-off is rather about maximizing the reduction in the headcount ratio or the poverty gap – the latter corresponding to a stronger focus on the poorest of the poor.

6 Conclusion

Using a global database covering 97% of the world's population, this paper shows that under assumptions of distribution-neutral growth, the World Bank's goal of achieving less than 3% extreme poverty by 2030, as well as the United Nations' Sustainable Development Goal of complete eradication of poverty, will be very difficult to reach by 2030. It also shows that these goals become more viable by reducing inequalities, for any given growth rate of the mean. Conversely, regressive distributional changes can severely limit the way in which growth contributes to poverty reduction.

Motivated by the Sustainable Development Goal 10 on inequality, we modeled inclusive growth in terms of lowering the Gini index in every country. The poverty impact of more inclusive growth defined in this way is different across countries and depends on the initial level of poverty, as well as the shape of the distribution, the precise growth incidence curve used and the growth rate. At high levels of initial poverty, reducing the Gini index could lead to a decrease in the pace of poverty reduction in the short term compared with a distribution-neutral growth scenario. In other words, for a country with a high headcount ratio, the welfare of the marginally poor may be growing slower when lowering the Gini than in a distribution-neutral scenario. This highlights a certain trade-off between focusing on the poor within every country and the poor according to an international poverty line. Nevertheless, in such cases the poorest of the poor still receive a growth premium and thus the poverty gap and severity are reduced.

While many of the findings are intuitive, one of the contributions of the paper is to quantify these effects using plausible distributional changes. A 1% annual decline in each country's Gini index is shown to have a bigger impact on global poverty than if each country experiences 1 pp higher annual growth rates than forecast. It is important to highlight that making growth more pro-poor as simulated in this paper does not impose a large cost on the rest of the distribution. Because of the large income share of the top of the distribution, the reduction in the growth rate of the wealthiest individuals necessary to ensure that the bottom grows substantially faster than the mean is relatively small. For example, in the case of Côte d'Ivoire, a convex growth incidence such that the Gini decreases by 1% means that households at the 10th percentile grow 2.5 pp faster than the mean, yet the growth at the 90th percentile is reduced by only 0.5 pp. In other words, the distributional changes simulated in this paper are not unrealistic, and as we have shown, making growth more pro-poor will be crucial for reaching the poverty goals set by the global development community.

7 References

- Atamanov, Aziz, Dean M. Jolliffe, Christoph Lakner and Espen Beer Prydz., "Purchasing Power Parities used in Global Poverty Measurement ", World Bank Group Global Poverty Monitoring Technical Note, no. 5., September 2018.
- Abdelkrim, A. and J.-Y. Duclos, "DASP: Distributive Analysis Stata Package," PEP, World Bank, UNDP and Université Laval, 2007
- Ahluwalia, M.S., "Income Inequality: Some Dimensions of the Problem." In H. Chenery, M.S. Ahluwalia, C.L.G. Bell, J.H. Duloy and R. Jolly (eds.), *"Redistribution with Growth,"* The World Bank, Oxford University Press, 1974.
- Birdsall, N., N. Lustig and C. J. Meyer, "The Strugglers: The New Poor in Latin America?," *World Development*, 60, 132-146, 2014.
- Breiman, L., Friedman, J., Stone, C., and Olshen, R. (1984). *Classification and Regression Trees.* Taylor & Francis, Belmont.
- Chandy, L., N. Ledlie, and V. Penciakova, "The final countdown: Prospects for ending extreme poverty by 2030," *Global Views Policy Paper 2013-04*, The Brookings Institution: Washington DC, 2013.
- Chen, S. and M. Ravallion, "The Developing World is Poorer than We Thought, But No Less Successful in the Fight Against Poverty," *The Quarterly Journal of Economics*, 125(4), 1577-1625, 2010.
- Deaton, Angus. "Measuring poverty in a growing world (or measuring growth in a poor world)." *Review of Economics and statistics* 87, no. 1 (2005): 1-19.
- Dhongde, S. and C. Minoiu, "Global Poverty Estimates: A Sensitivity Analysis," *World Development*, 44, 1-13.
- Edward, P. and A. Sumner, 2014: "Estimating the Scale and Geography of Global Poverty Now and in the Future: How Much Difference Do Method and Assumptions Make?" *World Development*, 58, 67-82, 2013.
- Ferreira, F. and P. Leite, "Policy Options for Meeting the Millennium Development Goals in Brazil: Can micro-simulations help?" *Economia Journal of the Latin American and Caribbean Economic Association*, vol. 0 (Spring 20), pages 235-280, January, 2003.
- Foster, J., J. Greer and E. Thorbecke, "A Class of Decomposable Poverty Measures," *Econometrica*, 52(3), 485-97, 1984.
- GMD (Global Monitoring Database), Global Solutions Group on Welfare Measurement and Capacity Building, Poverty and Equity Global Practice, World Bank, Washington, DC.
- Hellebrandt, T. and P. Mauro, "The Future of Worldwide Income Distribution," Working Paper Series 15-7, Peterson Institute for International Economics, 2015.
- Higgins, M. and J.G. Williamson "Explaining Inequality the World Round: Cohort Size, Kuznets Curves, and Openness," *Southeast Asian Studies*, 40(3), 2002.
- Hoy, Chris, and Emma Samman. 2015. "What If Growth Had Been as Good for the Poor as Everyone Else?" Report (May), Overseas Development Institute, London.
- Hillebrand, E., "The Global Distribution of Income in 2050," *World Development*, 36(5), 727-40, 2008.
- International Monetary Fund, "Fiscal Policy and Income Inequality," IMF Policy Paper. Washington, 2014.
- Jolliffe, D. and E. Prydz., "Global poverty goals and prices: how purchasing power parity matters," Policy Research Working Paper Series 7256, The World Bank: Washington DC, 2015.
- Jolliffe, D., and E. Prydz., "Estimating International Poverty Lines from Comparable National Thresholds." *The Journal of Economic Inequality* 14 (2): 185–98. 2016.
- Kakwani N., "Poverty and Economic Growth with Application to Côte d'Ivoire," *Review of Income and Wealth*, 39(2), 121-139, 1993.
- Karver, J., C. Kenny and A. Sumner, "MDGs 2.0: What Goals, Targets and Timeframe?," CGD Working Paper, Center for Global Development: Washington DC, 2012.
- Kraay, A. "When Is Growth Pro-Poor? Evidence from a Panel of Countries," *Journal of Development Economics*, 80(1), 198-227, 2006.

- Lakner, Christoph, Daniel Gerszon Mahler, Minh C. Nguyen, Joao Pedro Azevedo, Shaohua Chen, Dean M. Jolliffe, Espen Beer Prydz and Prem Sangraula., "Consumer Price Indices used in Global Poverty Measurement ", World Bank Group Global Poverty Monitoring Technical Note, no. 4., September 2018.
- Lakner, C., Negre, M. and Prydz, Espen Beer, 2014. "Twinning the Goals: How Can Promoting Shared Prosperity Help to Reduce Global Poverty?," Policy Research Working Paper Series 7106, The World Bank.
- Minoiu, C. and S. Reddy, "Kernel density estimation on grouped data: the case of poverty assessment," *Journal of Economic Inequality*, 12(2), 163-89, 2014.
- Ncube, M., Z. Brixiova and Z. Bicaba, "Can Dreams Come True? Eliminating Extreme Poverty in Africa by 2030," IZA Discussion Paper Series No. 8120, 2014:.
- Palma, J. G., "Homogeneous Middles vs. Heterogeneous Tails, and the End of the 'Inverted-U': It's All About the Share of the Rich," *Development and Change*, 42(1), 87-153, 2011.
- Pinkovskiy, Maxim, and Xavier Sala-i-Martin. "Lights, Camera... Income! Illuminating the national accounts-household surveys debate." *The Quarterly Journal of Economics* 131, no. 2 (2016): 579-631.
- PovcalNet: the online tool for poverty measurement developed by the Development Research Group of the World Bank, <http://iresearch.worldbank.org/PovcalNet>.
- Ravallion, M., "Growth, Inequality and Poverty: Looking Beyond Averages," *World Development*, 29(11), 1803-1815, 2001.
- Ravallion, M., "Measuring Aggregate Welfare in Developing Countries: How Well Do National Accounts and Surveys Agree?" *The Review of Economics and Statistics*, 85(3), 645-652, 2003.
- Ravallion, M., "A Global Perspective on Poverty in India," *Economic and Political Weekly*, 43(43), 33-37, 2008.
- Ravallion, M., "Do Poorer Countries have less Capacity for Redistribution?" *Journal of Globalization and Development*, De Gruyter, vol. 1(2), pages 1-31, December, 2009.
- Ravallion, M., "How Long Will It Take to Lift One Billion People Out of Poverty?" *The World Bank Research Observer*, 28(2) 2013.
- Ravallion, M. and S. Chen, "Measuring pro-poor growth," *Economic Letters*, 78(1), 93-99, 2003.
- Rodrik, D., "The Past, Present, and Future of Economic Growth," *Challenge* 57(3), 5-39, 2014.
- Shorrocks, A. and G. Wan., "Ungrouping Income Distributions," Working paper 2008/16, UNUWIDER, 2008.
- United Nations, "Open Working Group proposal for Sustainable Development Goals," United Nations: New York, 2014.
- Yoshida, N., H. Uematsu and C.E. Sobrado, "Is Extreme Poverty Going to End? An Analytical Framework to Evaluate Progress in Ending Extreme Poverty," Policy Research Working Paper 6740, The World Bank: Washington DC, 2014.
- World Bank, "Prosperity for All/Ending Extreme Poverty: A Note for the World Bank Group Spring Meetings 2014," Washington DC, 2014.
- World Bank, "A Measured Approach to Ending Poverty and Boosting Shared Prosperity: Concepts, Data, and the Twin Goals," Policy Research Report, The World Bank: Washington DC, 2015.
- World Bank, "Poverty and Shared Prosperity 2016: Taking on Inequality," The World Bank: Washington, DC, 2016.
- World Bank, "Poverty and Shared Prosperity 2018: Piecing Together the Poverty Puzzle", The World Bank: Washington DC, 2018a.
- World Bank, 2018b: Global Shared Prosperity Database. <https://datacatalog.worldbank.org/dataset/global-database-shared-prosperity>
- Zeileis, Achim, Torsten Hothorn, and Kurt Hornik. "Model-based recursive partitioning." *Journal of Computational and Graphical Statistics* 17, no. 2 (2008): 492-514.