



# How much does reducing inequality matter for global poverty?

Christoph Lakner<sup>1</sup> · Daniel Gerszon Mahler<sup>1</sup> · Mario Negre<sup>1,2</sup> · Espen Beer Prydz<sup>1,3</sup>

Received: 15 May 2019 / Accepted: 11 August 2021 / Published online: 2 March 2022

© World Bank under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2021

## Abstract

The goals of ending extreme poverty by 2030 and working towards a more equal distribution of incomes are part of the United Nations' Sustainable Development Goals. Using data from 166 countries comprising 97.5% of the world's population, we simulate scenarios for global poverty from 2019 to 2030 under various assumptions about growth and inequality. We use different assumptions about growth incidence curves to model changes in inequality, and rely on a machine-learning algorithm called model-based recursive partitioning 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 World Bank forecasts and historically observed growth rates, our simulations suggest that the number of extreme poor (living on less than \$1.90/day) will remain above 600 million in 2030, resulting in a global extreme poverty rate of 7.4%. If the Gini index in each country decreases by 1% per year, the global poverty rate could reduce to around 6.3% in 2030, equivalent to 89 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 1 percentage point above forecasts. We also study the impact of COVID-19 on poverty and find that the pandemic may have driven around 60 million people into extreme poverty in 2020. If the pandemic increased the Gini index by 2% in all countries, then more than 90 million may have been driven into extreme poverty in 2020.

**Keywords** Poverty · Inequality · Inclusive growth · COVID-19 · SDGs · Simulation · Machine-learning

---

✉ Christoph Lakner  
clakner@worldbank.org

<sup>1</sup> World Bank, Washington, DC, USA

<sup>2</sup> German Development Institute, Bonn, Germany

<sup>3</sup> Norwegian Institute of Public Health, Oslo, Norway

## 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, 2018). 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 150 countries and grouped data for an additional 16 countries, allowing for an unprecedented data coverage of 97.5% of the world’s population. *Second*, we model the impact of distributional changes on future trajectories of global poverty by changing a country’s 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 7.4% in 2030 if growth is distribution-neutral and follows World Bank forecasts until 2021 and country-specific historical growth averages from then on. Under a scenario in which the Gini index of each country decreases by 1% per year, the global poverty rate falls to 6.3% – equivalent to 89 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 above World Bank forecasts. Even under the most optimistic scenarios we consider – where the Gini decreases 2% annually and the annual growth rate exceeds World Bank forecasts and historical averages by 2 percentage points – the poverty rate in Sub-Saharan Africa would remain around 20% in 2030 and the global target of 3% would not be met.

We also study how COVID-19 affects these projections of global poverty. Our baseline scenario suggests that the pandemic has driven 60 million people into extreme poverty in 2020. If all countries’ Gini indices increased by 2% in 2020 due to the pandemic, then 94 million will have been driven into extreme poverty. This is a larger effect than if all countries growth forecasts are 2 percentage points lower than anticipated in which case the

pandemic is expected to drive 82 million people into extreme poverty. Hence, our finding that percentage changes in the Gini matter more than percentage point changes in growth carries over to the estimated COVID-19 impacts. This finding, namely that a 1% decline in the Gini index implies greater poverty reduction than a 1 percentage point increase in growth, of course does not imply that it will be easier to implement inequality-reducing than growth-enhancing policies in practice. Political economy reasons might well stand in the way of sustained decreases in inequality.

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 muted effect on *global* inequality, for which between-country differences matter greatly (Anand & Segal, 2008; Lakner & Milanovic, 2016). 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 eliminate poverty. Our baseline results are based on a linear growth incidence curve (GIC), but in a robustness check we use a convex 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 index in all countries has about the same impact on global poverty as a 2 percentage point 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 imposed distribution-neutral growth, thus ignoring any future changes in within-country inequality (Birdsall et al., 2014; Karver et al., 2012; Hellebrandt & Mauro, 2015). Others have projected distribution-neutral growth but chosen initial distributions with different levels of inequality (Ravallion, 2013; Edward & Sumner, 2014). Other studies simulate additional distributional changes by extrapolating the trend in the Q5/Q1 ratio (Edward & Sumner, 2014; Hillebrand, 2008; Higgins & 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).

Two main approaches are used in the literature to project poverty forward which can produce quite different results (Dhongde & Minoiu, 2013; Edward & 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 by projecting historical or forecasted growth rates into the future (Karver et al., 2012; Chandy et al., 2013). It should be noted that some of these methods were primarily used as an attempt to set a goal for global poverty in 2030 rather than as an attempt to forecast the poverty rate in 2030. For the former, one may choose a method that sets an ambitious yet achievable goal, rather than a method that results in an estimate one expects to be realized. We base our projections on forecasted growth rates from the World Bank's Global Economic Prospects (GEP) until 2021 (the last year for which growth forecasts were available at the time of writing), and country-specific historical growth rates (from 2008 to 2018) to project forward

from 2022 to 2030. Our projections adjust for differences between growth from household surveys and national accounts.

Several other papers have estimated the impact of COVID-19 on global poverty. In particular, Sumner et al. (2020) explore what happens if all countries' growth rates decline a fixed amount while Laborde et al. (2020) estimate the impact using a general equilibrium model. We estimate the impact using household survey data and two vintages of growth projections for 166 countries, allowing us to compare COVID-19 poverty projections with counterfactual projections from just before COVID-19.

We model distributional changes and growth rates in GDP independently of each other. Although the 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.

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, poverty measures, and passthrough rates. Section 6 discusses the reasons why poverty reduction might be slowing down while Section 7 concludes.

## 2 Conceptual framework

In this paper, we model how changes in each country's Gini index impact poverty towards 2030. 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 focus on two functional forms of the GIC.<sup>1</sup> 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

$$y_i^* = y_i(1 + g_i) \quad (1)$$

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 implement changes in the Gini index 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  $\gamma$  together with a tax and transfer scheme which taxes everyone at a rate  $\tau$  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  $p$ , 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

<sup>1</sup> 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.

$Gini(y)^* = (1 - \tau)Gini(y)$ . In other words, the tax rate imposed,  $\tau$ , is equivalent to the percentage change in Gini observed,  $\alpha$ , such that  $\tau = -\alpha$ . We can express the final incomes as a function of the initial income, mean income, and changes in the Gini index:

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

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

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

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. This GIC is concave when inequality is increasing ( $\tau < 0$ ). For simplicity, we refer to this functional form of the GIC as convex throughout the paper regardless of whether the Gini index increases or decreases.

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

$$g_i = \delta - \theta p_i \tag{4}$$

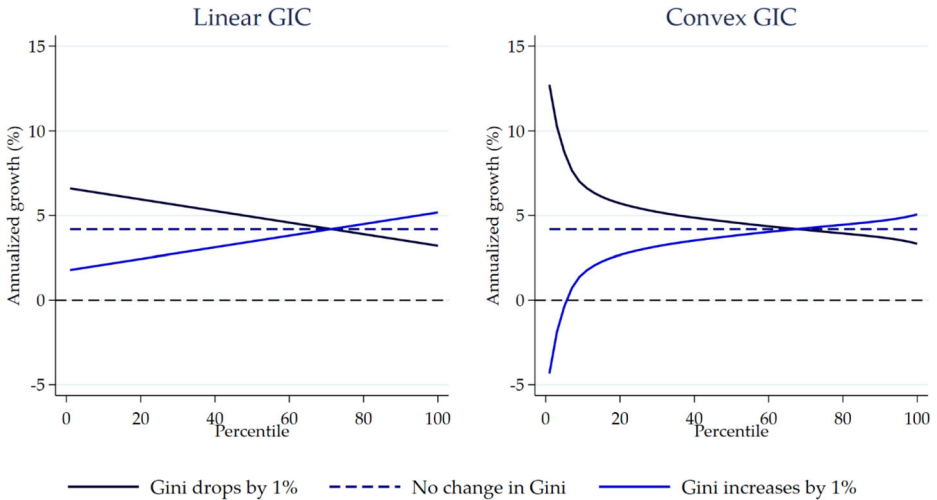
Substituting (4) into (1), we can obtain the following expression for the income of percentile group  $i$  in the final period:

$$y_i^* = (1 + \delta)y_i - \theta y_i p_i \tag{5}$$

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  $\theta$  and the tax increases proportionally with the rank – combined with a transfer where every person receives a share  $\delta$  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,  $\alpha$ , and the parameters of the linear GIC.

To illustrate how the convex and linear GICs could look like in practice, we use the welfare distribution for Côte d’Ivoire from 2018 from PovcalNet. From 2018 to 2019 the World Bank’s Global Economic Prospects (GEP) suggest that real GDP per capita in Côte d’Ivoire grew by 4.3%. Figure 1 explores how this growth can be distributed if inequality stays unchanged or if the Gini index increases or decreases by 1% (ignoring for the moment whether 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 index by 1% does not have to impose a large cost (in terms of reduced growth) on the top of the distribution. Because of the larger income share of the top of the distribution, the reduction in the growth rate of the richest 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 percentage points faster than the mean, yet only reduces the growth at the 90th percentile by 0.5 percentage point.



**Fig. 1** Different growth incidence curves compatible with the same change in the Gini index. Note: Growth incidence curves (GICs) drawn using data for Côte d’Ivoire from PovcalNet for 2018 under different assumptions about how much inequality changes, and in what manner inequality changes. The mean is assumed to grow at 4.3 %, according to data from the GEP.

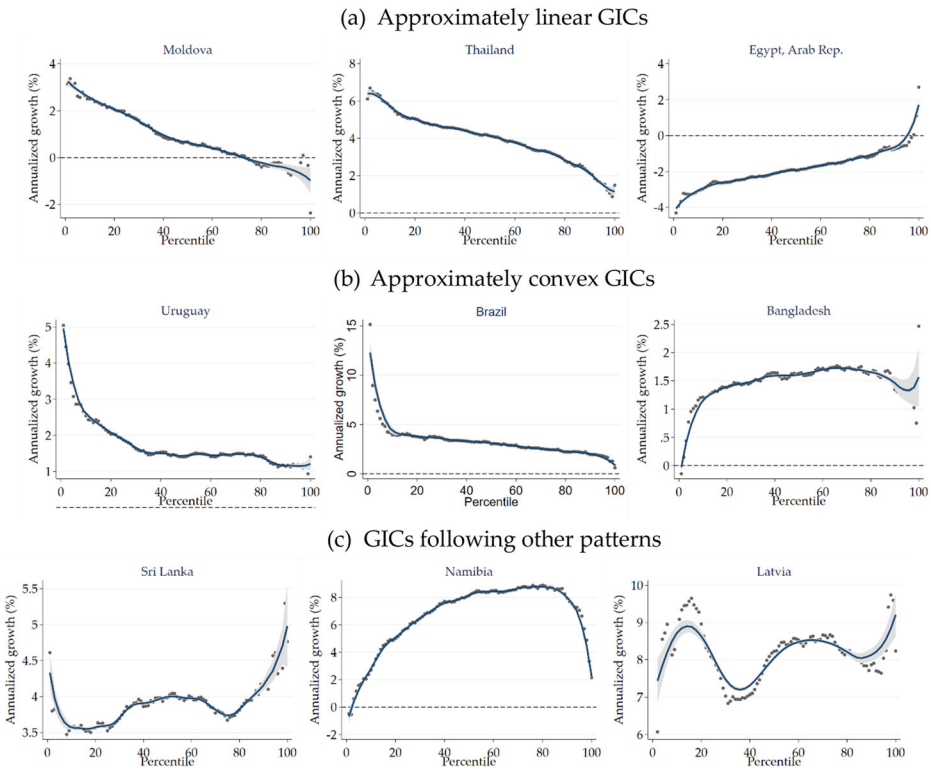
In our baseline simulations we use linear GICs for three reasons: First, it is probably the simplest realistic pro-poor GIC that can be constructed. Second, it constitutes a relatively conservative pro-poor distributional change, in contrast to the convex GIC, which may provide a too optimistic picture 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 index 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.

Nonetheless, the convex GIC has some advantages: First, it intuitively relates to public policy, as it represents the outcome of a simple tax and transfer scheme. 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. In contrast, for the linear GIC, we are forced to use an algorithm that iteratively changes the slope of the GIC until it matches the desired  $\alpha$ . Third, it is directly linked to differences in growth rates of the bottom 40% and the overall distribution, also called shared prosperity, which is the first SDG target on inequality.<sup>2</sup> 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 2020), which provides a list of

<sup>2</sup> The percentage change in the Gini index,  $\alpha$ , and the shared prosperity premium,  $m$  (the difference in growth of the bottom 40% and the mean), are related as follows:  $\alpha = \frac{m}{(1+\gamma) \left( \frac{0.4}{s_{40}} - 1 \right)}$ , where  $s_{40}$  is the income share of the

bottom 40%. For a given income share of the bottom 40% and overall growth rate, there is a linear relationship between the size of the tax rate, the percentage change in the Gini index, and the shared prosperity premium. For more details on the relationship between the convex GIC and shared prosperity see the appendix of Lakner et al. (2014).



**Fig. 2** Empirically observed growth incidence curves. Note: Empirically observed growth incidence curves using the surveys in World Bank (2020).

259 spells with a 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.

An alternative to using a theoretically defined GIC would be to impose one that has been observed in practice 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 challenging for the many countries that lack comparable data over time.

### 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.<sup>3</sup> Most of the

<sup>3</sup> Data from PovcalNet can be accessed at <http://iresearch.worldbank.org/PovcalNet/home.aspx> or directly through Stata or R (Castaneda et al., 2019a). This paper uses PovcalNet data from the March 2020 vintage.



data in PovcalNet comes from the Global Monitoring Database, which is the World Bank's repository of multitopic income and expenditure household surveys used to monitor global poverty. PovcalNet contains more than 1900 surveys from 166 countries covering 97.5% 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 150 of the countries, housing 69% of the world's population, micro data are available. For an additional 8 economies (Australia, Canada, Germany, Israel, Japan, South Korea, Taiwan, and the United States), or 9% of the world's population, grouped data of 400 bins are available. For the purposes of these projections, we treat these bins as microdata. Finally, for China and seven other countries constituting about 19% of the world's population, only decile or ventile shares and the overall mean are available. Aside from China, this concerns Algeria, Guyana, Suriname, Turkmenistan, Trinidad & Tobago, Venezuela, and the United Arab Emirates. For these countries, we follow PovcalNet and fit a General Quadratic Lorenz curve and a Beta Lorenz curve, choosing the one that gives the best fit, and use it to recover a full distribution.<sup>4</sup>

### 3.2 Growth scenarios

Our starting point in each country is the welfare distribution the World Bank uses to measure poverty for the country in 2018, which is the latest year with poverty estimates at the time of writing. These welfare distributions are based on (often) extrapolated distributions from household surveys. The median year of data for these estimates is 2016, but the range spans from 1992 to 2018.<sup>5</sup>

To project poverty forward, a commonly used strategy is to rely on historical annualized growth rates. COVID-19 makes this a quite unattractive option due to the high likelihood of an increase in poverty in 2020. From 2018 to 2021, we therefore use the growth projections from the June 2020 edition of the World Bank's Global Economic Prospects (GEP) to account for the impact of COVID-19 on economic activity.<sup>6</sup> 2021 is the last year for which growth projections are available. Beyond that, one could use the annualized growth in the forecasting period, or the last growth rate of the forecasting period to project forward towards 2030. COVID-19 makes these options unattractive as well due to the extreme growth rates observed in both 2020 and 2021. Therefore, beyond 2021 we use three different scenarios based on historical growth rates: each country grows according to its annualized growth rate from national accounts for the last 5, 10, or 20 years leading up to 2018 (1998–2018, 2008–2018, 2013–2018). The simulations relying on the 20-year historic growth rates may be optimistic, as

<sup>4</sup> 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.

<sup>5</sup> If countries do not have survey data for 2018, PovcalNet extrapolates their latest survey to 2018 using growth in GDP per capita or Household Final Consumption Expenditure per capita assuming distribution-neutrality (Prydz et al., 2019). The only country for which an extrapolated estimate is not available in 2018 is India. For India, we follow the extrapolation approach used for the other countries to generate an estimate of the distribution in 2018.

<sup>6</sup> For the economies not in GEP, we use growth forecasts from IMF's World Economic Outlook. Syria does not have growth projections in either of these sources. In this case, we use the regional average growth forecast of the Middle East and North Africa region to project forward.



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.

Our preferred source of historical growth data is growth in real GDP per capita from national accounts as reported in the World Development Indicators (WDI), which includes data up to 2018 at the time of writing. When such data are not available for the whole period, we complement it with growth data used by PovcalNet for monitoring global poverty. Most of the added sources are from the Maddison Project Database (Prydz et al., 2019).

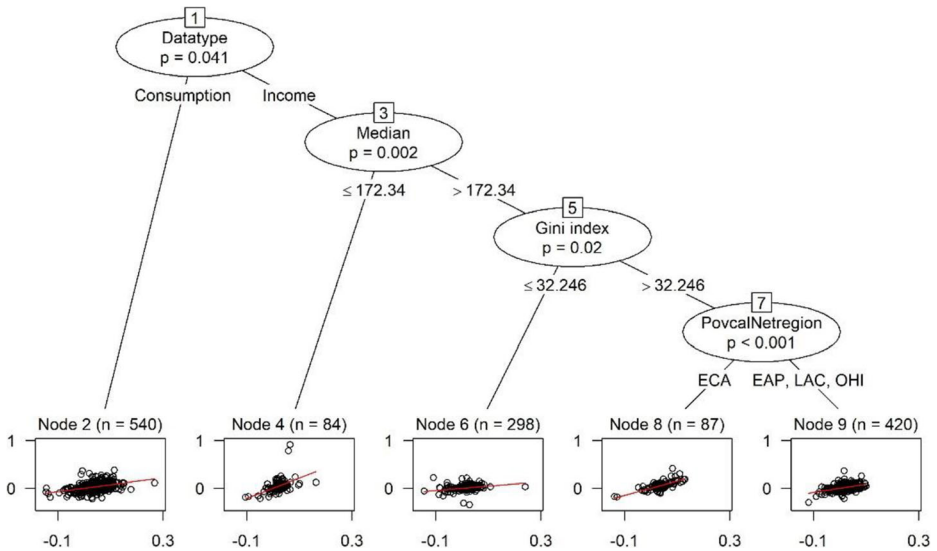
### 3.3 The relationship between growth in national accounts and surveys

A challenge with using growth rates in GDP per capita to project poverty forward is that prior evidence has shown that only a fraction of growth observed in national accounts is passed through to 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 annualized growth in survey means on annualized growth in real GDP per capita, under the constraint that the intercept is zero,  $g_{survey} = \beta * g_{GDPpc} + \epsilon$ , and use  $\beta$  as the fraction of growth in GDP per capita that is passed-through to welfare observed in surveys. Using 1429 spells with comparable household survey data suggests that  $\beta = 0.85$ . Each spell relies on two adjacent comparable surveys from the same country with welfare measured in the same way, either income or consumption (World Bank, 2019).

Yet there is no reason to believe that  $\beta$  is constant across different contexts. It may differ by geographical region, by income level, by whether income or consumption is used, 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 & Ravallion, 2010; Chandy et al., 2013; and Corral et al., 2020), but it is not entirely clear on what basis to select the variables to be included.

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 takes 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 per capita, and the year of the survey. The algorithm is a variant of classification and regression trees, pioneered by Breiman et al. (1984), and works in the following manner:

1. Run the regression  $g_{survey} = \beta * g_{GDPpc} + \epsilon$  on all relevant data.
2. Add interactions between  $g_{GDPpc}$  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



**Fig. 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.

variable. If the lowest p-value is greater than 0.05, no split is made, and the algorithm stops (suggesting that there is no evidence in favor of passthrough rates differing).

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 to split the sample into two. Splits are only made if at least 10 observations will be in each subsample.
5. The algorithm is repeated from the beginning by applying it to observations in each of the two subsamples separately.

Figure 3 and Table 1 show the results of model-based recursive partitioning using our data at hand. There is significant evidence in favor of the data type mattering for passthrough rates. Observations using income have a passthrough rate of 1.01, while observations using consumption have a passthrough rate of 0.72. With a p-value of 0.041, 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 median matters for determining the passthrough rate. Cases with a median less than 172 USD per person per month in 2011 PPPs (or 5.7 per day) have a passthrough rate of 2.11 while observations with a median above this threshold have a passthrough rate of 0.87, and so forth. Table 1 contains more details on the Wald tests, the splits conducted, and the associated passthrough rates. Two-thirds of all cases are predicted to have a passthrough rate between 0.72 and 0.86.

Using model-based recursive partitioning is only one machine-learning method amongst many that endogenize the interactions to include. We find this method attractive because it is

**Table 1** Details on decision tree algorithm

Node	Obs.	$\beta$	p-values from Wald tests										
			Data-type	Gini	Median	Mean	GDP	World Bank region	PovcalNet region	Year	Population	Head-count	
1	1429	0.85	<b>0.04</b>	0.54	1.00	0.86	0.08	0.98		0.91	0.99	0.97	0.99
2	540	0.72	----	1.00	0.96	0.92	0.52	1.00		1.00	0.99	0.79	0.99
3	889	1.01	----	0.27	<b>0.00</b>	0.14	0.59	1.00		0.59	0.54	0.98	0.02
4	84	2.11	----	1.00	0.84	1.00	0.94	0.00		0.00	0.87	1.00	0.98
5	805	0.87	----	<b>0.02</b>	0.82	0.71	0.91	1.00		0.70	0.63	0.81	0.06
6	298	0.46	----	0.86	1.00	1.00	1.00	1.00		0.90	1.00	0.93	0.36
7	507	1.08	----	0.11	1.00	1.00	0.37	0.18		<b>0.00</b>	0.63	0.70	1.00
8	87	1.51	----	0.99	1.00	1.00	1.00	----		----	0.99	1.00	1.00
9	420	0.86	----	1.00	0.91	1.00	1.00	1.00		1.00	1.00	1.00	0.98

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. “—” indicates that no test can be conducted since there is no variation in the input variable in question in the particular subsample. In node 4 the region variables are significant but no splits are made since the desired splits would leave less than 10 observations in one of the subsamples.

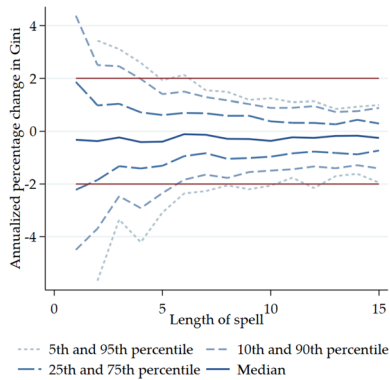
specifically designed to test whether a parameter of interest differs by subgroups, it relies on statistical tests, and it is easy to visualize. A shortcoming of this method is that its coarseness means that small changes in the underlying data could change the predictions. In Section 5.2 we discuss our choice in more detail, show robustness checks using the lasso and a constant passthrough rate across all observations, and compare the out-of-sample performance.

### 3.4 Inequality scenarios

We consider five different scenarios for changes in the Gini index; that it changes by -2%, -1%, 0%, 1% and 2% per year beginning in 2019. If a country starts with a Gini index of 0.40 in 2019 (which is close to the median Gini using the latest survey for each country), under our five different scenarios, it would end up with a Gini of 0.32, 0.36, 0.40, 0.45 and 0.50 in 2030, respectively.

Evaluating the plausibility of these Gini changes requires comparable data across countries over time. Using the comparability database associated with PovcalNet (World Bank, 2019), we can recover 8,322 comparable spells. Figure 4 shows the annualized percentage change in the Gini index from these spells, as a function of spell length, where each spell has been given a weight equal to the inverse of the number of spells by country-spell length.<sup>7</sup> This means that for a given spell length, each country gets a weight of 1 irrespective of how many spells can be observed for this particular length.

<sup>7</sup> For the passthrough rate analysis in the previous section, we recovered many fewer comparable spells (1429) since we only looked at adjacent surveys for a particular country. Here we also consider surveys that are comparable even if they are not adjacent (meaning other surveys were carried out in between). We are applying weights in order to get a balanced sample of countries at each spell length, to the extent possible. Still, since most countries do not have comparable surveys corresponding to all spell lengths, the set of countries at each spell length varies.



**Fig. 4** Observed annualized changes in the Gini. Note: Distribution of observed annualized changes in the Gini index using the 8,322 comparable spells from the surveys available in PovcalNet. Each spell is weighted by the inverse of the number of spells by country-spell length.

The figure reveals that Gini changes tend to be smaller the longer the spell length, suggesting that large changes in the Gini are difficult to sustain over long periods of time. For spell lengths of 11 years, which are equivalent to the 2019–2030 spell length we look at in this paper, annual declines of 1% per year are just below the 75th percentile while annual declines of 2% per year are around the 95th percentile of the distribution of changes in the Gini index. Thus, both of these seem plausible in a historical perspective. An annual increase in the Gini index of 1% is around the 5th percentile and is therefore also plausible. Annualized increases of 2%, however, have not been seen over an 11-year period.

### 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).<sup>8</sup>

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 data used by PovcalNet. Most CPIs are from the IMF's International Financial Statistics, while most PPP exchange rates are from the International Comparison Program (PPPs for household final consumption expenditure).<sup>9</sup> 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

<sup>8</sup> The povsim ado can be found at <https://github.com/danielmahler/InequalitySensitiveProjections>.

<sup>9</sup> We use the original 2011 PPPs as published in December 2014. Revised 2011 PPPs were published in May 2020. Atamanov et al. (2020) show that the impact of the PPP revisions on the global poverty estimates is very small.

poverty estimates across regions and globally. We use country-level population projections from the World Bank.<sup>10</sup> Finally, to arrive at regional and global poverty rates, we also need estimates for the 2.5% 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 poverty nowcasts to 2020 in an attempt to quantify the impact of COVID-19 on global poverty, and the relevance of assumptions about inequality and growth for quantifying this impact. Second, we project poverty towards 2030, both at the global and regional level, and explore what would happen if growth or inequality changes in a positive or negative direction. Unless otherwise specified, we focus on the international poverty line at \$1.90 per person per day in 2011 PPPs.<sup>11</sup>

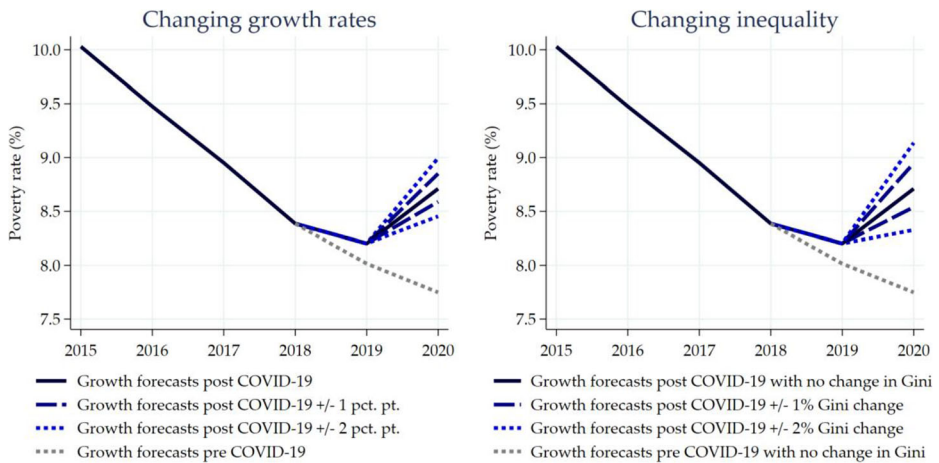
### 4.1 Nowcasting poverty: The impact of COVID-19, growth and inequality

Figure 5 shows nowcasts of poverty for 2020 (as well as projections to 2021) utilizing the growth forecasts from the June 2020 edition of the World Bank's GEP. As the crisis is still unfolding at the time of writing (June 2020), there is considerable uncertainty with regards to the growth impact and the impact of the pandemic on within-country inequality. Therefore, Fig. 5 also displays scenarios where the growth forecasts differ by -2, -1, 1 or 2 percentage points (compared to the GEP baseline) as well as scenarios where the Gini coefficient changes by -2, -1, 1 or 2% (using a linear GIC). In order to quantify the impact of the pandemic on global poverty, we compare these projections with the projections we would obtain using growth forecasts from the World Bank's GEP published in January 2020, which predates the global spread of COVID-19. Of course, other factors may have also worsened (or improved) countries' growth outlooks over these six months but it is safe to say that most of the changes in the forecasts are due to COVID-19.

With the new distribution-neutral forecasts, global poverty is projected to increase from 8.2% in 2019 to 8.7% in 2020, or from 630 million people to 675 million people. Compare this with the projected decline from 8.0 to 7.7% over the same time period using the previous GEP forecasts. The slight change from 8.2 to 8.0% for 2019 happens because the new vintage of the GEP also revised 2019 growth rates for some countries. Taking this into account, it means that COVID-19 is driving a change in our 2020 estimate of the global poverty rate of about 0.8 percentage points —  $(8.7\% - 8.2\%) - (7.7\% - 8.0\%)$ . Another way to put this is that the estimates suggest that COVID-19 will push 60 million people into extreme poverty in 2020, or equivalently that the number of

<sup>10</sup> These are available at <https://datacatalog.worldbank.org/dataset/population-estimates-and-projections>.

<sup>11</sup> See Ferreira et al. (2016) for a description of how the \$1.90 international poverty line has been derived.



**Fig. 5** Impact of COVID-19 on global poverty under different growth and Gini scenarios. Note: Projected global poverty rate measured at \$1.90 per person per day in 2011 PPPs assuming countries exceed or fall behind the growth projections from the GEP by 1 or 2 percentage points (left panel), or follow the GEP projections exactly but reduce/increase their Gini index by 1 or 2% (right panel). Inequality scenarios are based on linear growth incidence curves.

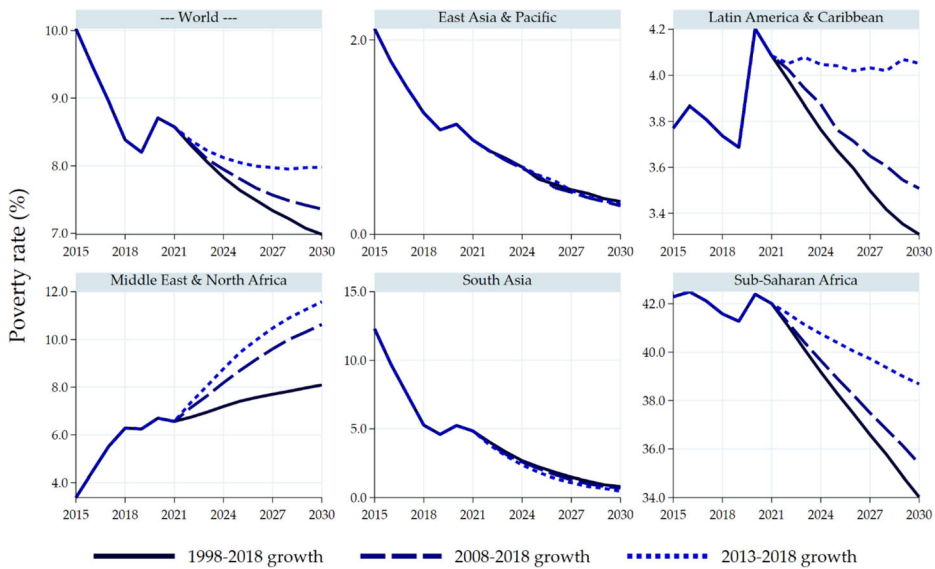
extreme poor increases by 10%. This marks the first time since the Asian Financial Crisis of 1997–1998 that the global poverty rate is increasing.

If growth in 2020 in all countries is 2 percentage points lower than GEP projections, COVID-19 would increase global poverty by 1.1 percentage points and the number of poor would increase by 82 million. If inequality increases by 2% in 2020 in all countries, then global poverty would increase by 1.2 percentage points and the number of poor would increase by 94 million.<sup>12</sup> The latter scenario would imply that the global progress on ending extreme poverty would be set back by three years. On the other hand, if all countries decrease their Gini coefficient by 2% in 2020, which could happen if countries successfully implement and expand social protection programs, then the number of people pushed into extreme poverty due to COVID-19 would be cut in half from the baseline, to around 30 million. Taken together, these results suggest that a given percentage change in inequality matters more for global poverty than a similar percentage point change in growth rates.

#### 4.2 Global and regional trajectories to 2030

Turning to poverty projections to 2030, we are faced with the problem that growth forecasts end in 2021. Figure 6 presents our simulated trajectories for the global poverty rate to 2030 for three different distribution-neutral growth scenarios: that countries beyond 2021 follow their growth patterns of the past 5, 10, or 20 years.

<sup>12</sup> Earlier we showed that annual increases of 2% in the Gini index do not persist for long periods. Yet, such changes are certainly observed from year to year.

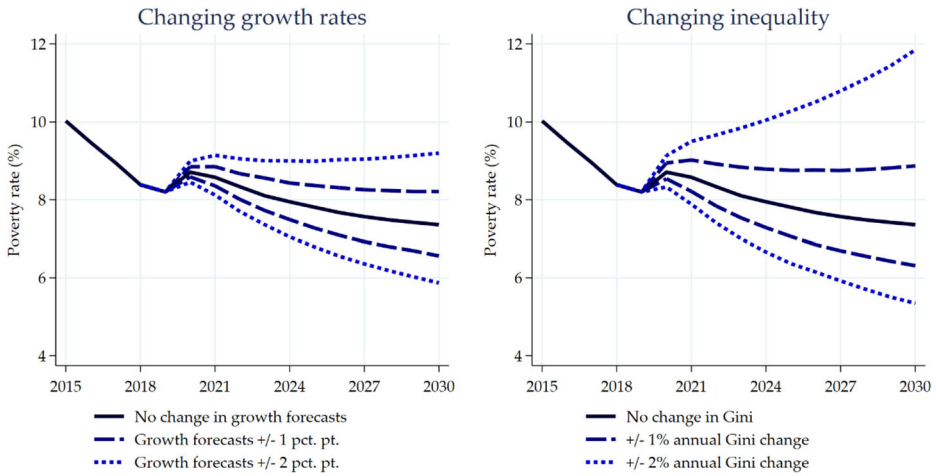


**Fig. 6** Global and regional distribution-neutral poverty projections to 2030. Note: Projected global and regional poverty rates measured at \$1.90 per person per day in 2011 PPPs assuming distribution-neutrality under three different growth scenarios (after 2021): countries follow their growth patterns of the past 5, 10, or 20 years. The regional definitions follow the World Bank’s classification, excluding the North America region and the Europe & Central Asia region where the share living below \$1.90 is small.

All scenarios put the global poverty rate in 2030 in the range of 7–8%. The scenario using historical growth rates from 1998 to 2018 is slightly more optimistic for some regions due to the high growth rates at the turn of the century. It is important to stress that these projections are not a prediction of what poverty will look like in 2030. Rather, they represent a hypothetical scenario of what would happen if all countries from 2022 onwards grow in accordance to what has occurred in the past. In Latin America & the Caribbean, using growth rates from the past 5 years results largely in a stagnation of poverty, while the other two scenarios decrease poverty substantially towards 2030. In the Middle East & North Africa, all scenarios yield increasing poverty rates towards 2030. The global poverty rate is largely driven by Sub-Saharan Africa, which in all three scenarios has poverty rates above 30% in 2030, while the other regions of the world have rates below 15% (the vertical axes differ across regions in Fig. 6).

Next, we look at how changing inequality or the growth rates impact global poverty. We focus on the scenario that uses annualized growth rates from 2008 to 2018 beyond 2021, and simulate the change in poverty if each country’s annual growth rate is 1 or 2 percentage points higher than its historical rate. We use the 2008–2018 scenario as our baseline since it is the intermediate scenario in terms of how optimistic it is for the future. 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 Fig. 7.

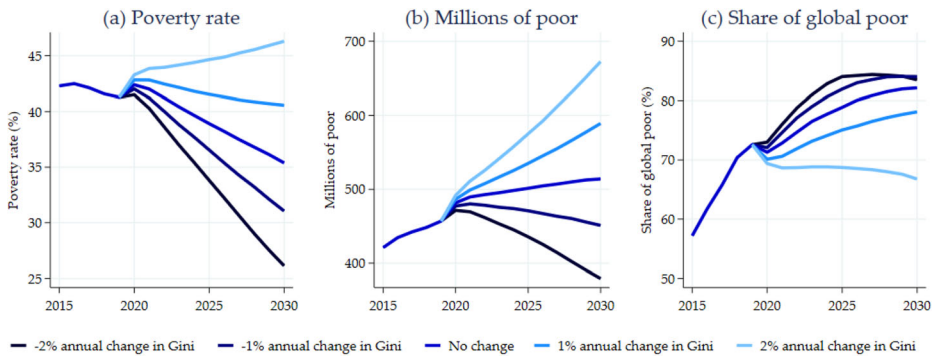




**Fig. 7** Simulations of global poverty under different growth and Gini scenarios. Note: Projected global poverty rate measured at \$1.90 per person per day in 2011 PPPs assuming that countries exceed or fall behind the growth projections from the GEP (until 2021) and 2008–2018 historical growth rates (from 2022) by 1 or 2 percentage points annually (left panel), or follow the growth rates exactly but reduce or increase their Gini index by 1 or 2% annually (right panel).

Decreasing the Gini index by 1% annually in each country has a larger impact on poverty than increasing growth 1 percentage point 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 6 or 9% with 1% annual decreases or increases in the Gini index, respectively. This does not speak to whether reducing inequality every year in countries is politically feasible, only that doing so would matter more for reducing global poverty than boosting growth rates (when comparing percentage changes in the Gini with percentage point changes in growth rates).

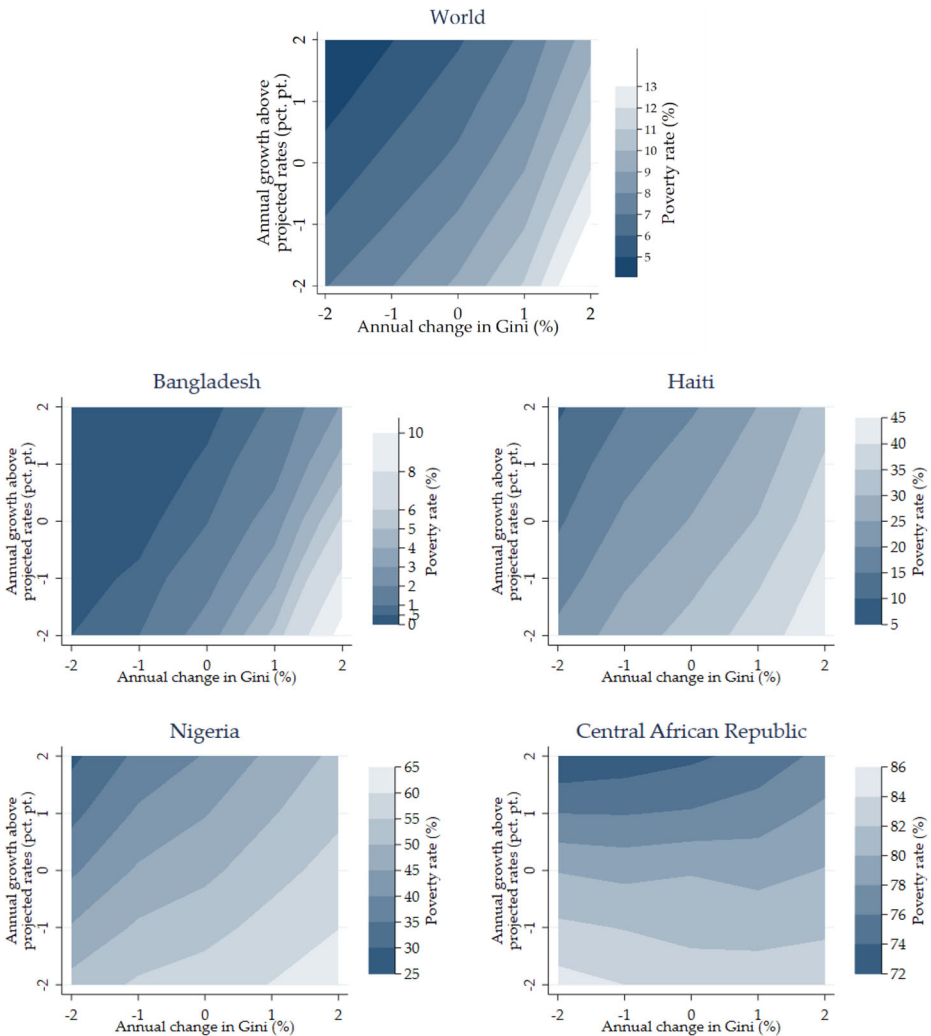
Changes in the Gini index are particularly relevant for Sub-Saharan Africa, where the poverty rate fluctuates between 30 and 40% with 1% annual decreases or increases in the Gini



**Fig. 8** Poverty projections in Sub-Saharan Africa. Note: Projected poverty rates in Sub-Saharan Africa measured at \$1.90 per person per day in 2011 PPPs assuming that countries follow the growth projections from the GEP and 2008–2018 historic growth, under five different scenarios about how inequality will change in each country.

index (Fig. 8a). Due to rapid projected population growth, only the scenarios that lower inequality are expected to decrease the number of poor in Sub-Saharan Africa (Fig. 8b). Since the inequality-reducing scenarios rapidly reduce poverty in other regions (with the exception of the Middle East and North Africa), the share of the global poor that live in Sub-Saharan Africa increases under these scenarios. Around 85% of the global poor would reside in Sub-Saharan Africa by 2030 if all countries experience a fall in inequality (Fig. 8c).

The combinations of scenarios changing the Gini index and making the growth rate higher or lower than GEP projections allows for the creation of iso-poverty curves. These curves, introduced by Ferreira and Leite (2003), show combinations of inequality changes and growth



**Fig. 9** Global and country-specific iso-poverty curves, 2030. Note: The figure shows different combinations of changes in the Gini index and in growth scenarios that result in the same poverty rate globally and for four selected countries. The flatter the curves, the more growth matters relative to reducing inequality.

changes resulting in the same level of poverty, as shown in Fig. 9. The flatness of the curves illustrates the relative role of growth and inequality in shaping poverty rates.

At a global level, the curves are steeper than 45 degrees, suggesting that reducing the Gini index by 1% with a linear GIC is more impactful than exceeding growth forecasts with a 1 percentage point. 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 Central African Republic, for example, in certain scenarios increasing the Gini lowers poverty.

These conclusions may depend on how we have set-up the simulations. If we used higher poverty lines, other countries would present a similar pattern to that seen in the Central African Republic. 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 will unambiguously reduce poverty. The conclusion that inequality is more important than growth may also be influenced by our choice of GIC. In the next section we will explore the robustness of the results to the choice of alternative poverty lines, poverty measures, and GICs. We will also look at how sensitive our projections are to our passthrough rate calculations.

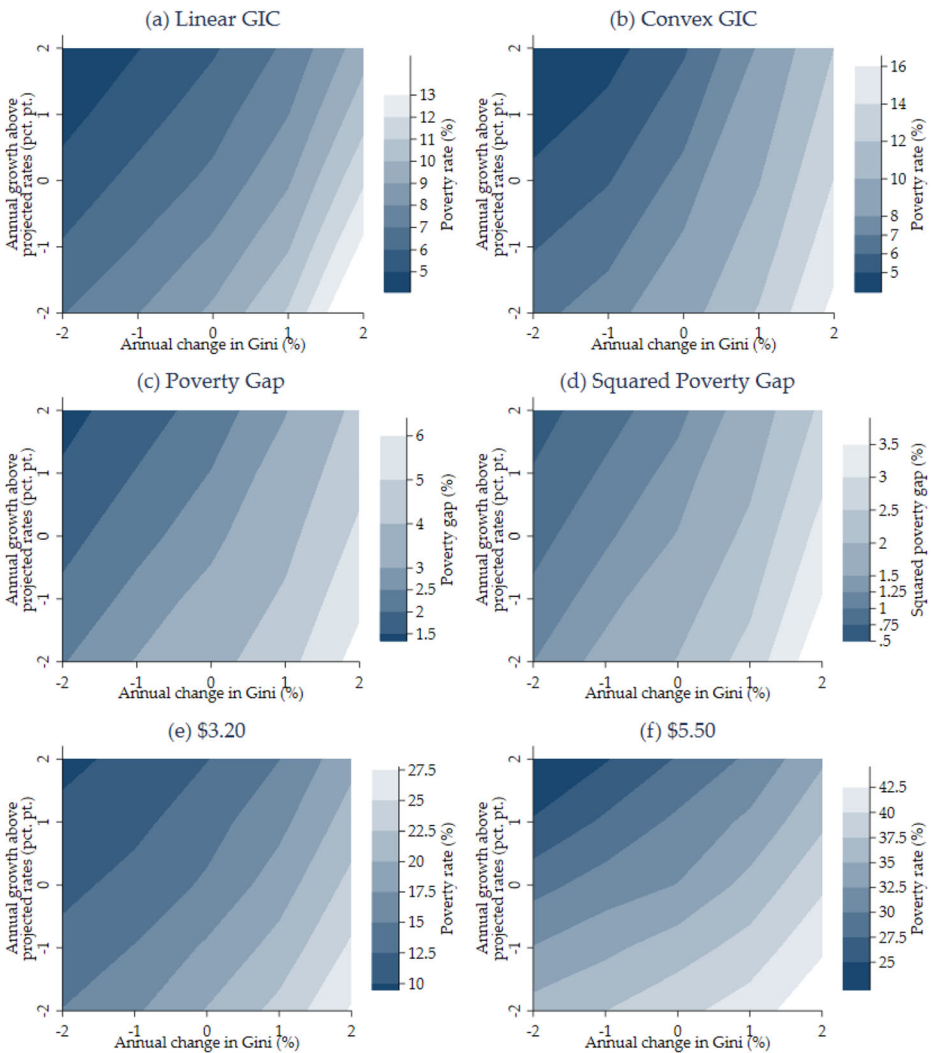
## 5 Robustness checks

### 5.1 Poverty measure, poverty line and growth incidence curve

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, the bottom of the distribution experiences larger 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 Fig. 10. 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 generally becoming steeper. Now a 1% annual reduction in the Gini matters as much as exceeding growth forecasts by 2 percentage points annually, as both bring the global poverty rate in 2030 to about 6%.

Next, 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. Distributional differences among the poor may be particularly important 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 iso-poverty curve become slightly steeper, meaning that changes in the Gini have an even larger impact on poverty reduction, relative to higher growth (panel c and d).

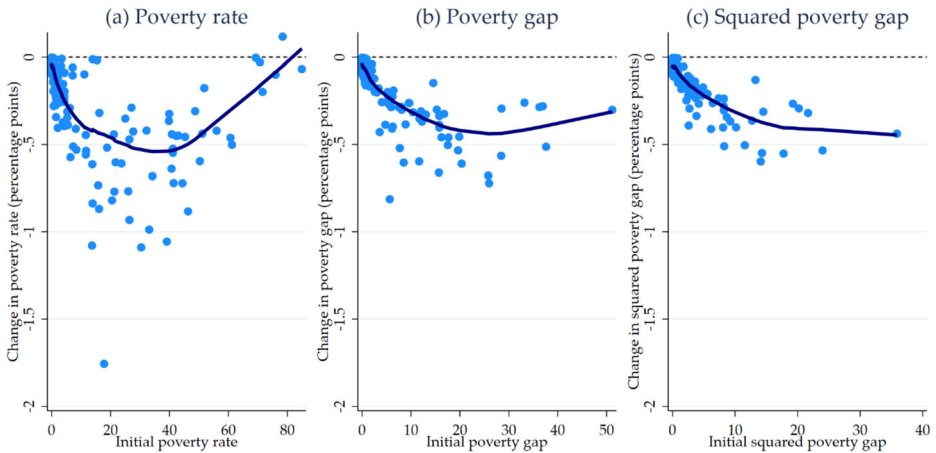
Finally, we use higher poverty lines. Specifically, we use the poverty lines of \$3.20 and \$5.50, which are official higher poverty lines of the World Bank (panel e and f). These lines are constructed to reflect typical national poverty lines in lower- and upper-middle income



**Fig. 10** 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 different poverty measures at the \$1.90 poverty line, the poverty gap and the squared poverty gap, respectively. Panel e and f show the poverty rate at higher poverty lines.

countries, respectively (Jolliffe & Prydz, 2016). With the \$3.20 line, a 1% annual decline in the Gini still has a larger impact on global poverty than exceeding growth forecasts by 1 percentage point per year, while at the \$5.50 line, they are about equally important for reducing poverty. This is despite the fact that almost half of the world lived below \$5.50 in 2015 (World Bank, 2018).

The impact of reducing the Gini index varies with the initial poverty rate, the shape of the GIC, the income distribution, and the measure of poverty. Figure 11 plots the impact of reducing the Gini index by 1% with a linear GIC for all countries, using their 2019 distributions as the starting point (without changing the mean).



**Fig. 11** Impact of reducing Gini by 1%, by poverty measure and poverty level. Note: The figure shows the percentage point change in poverty measures assuming the Gini decreases by 1% with 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 tradeoff between decreasing the poverty rate and decreasing inequality for very poor countries. Panels b and c show that reducing inequality unambiguously decreases both the poverty gap and the squared poverty gap even for high initial headcount ratios. This indicates that the tradeoff 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.

## 5.2 Passthrough rates

Our results were based on using model-based recursive partitioning to estimate the fraction of growth in GDP per capita that is passed through to growth in welfare observed in household surveys. This method is only one method amongst many to estimate passthrough rates. We employed this method because it is designed to test whether a parameter of interest – here the passthrough rate – differs by subgroups. That said, other machine-learning methods could also be adapted to this setting by including interactions between GDP per capita and relevant variables. In contrast to many other machine-learning approaches, model-based recursive partitioning has the advantage of being quite transparent in the sense that it is a sequence of statistical tests which can be visualized. Its coarseness on the other hand, implies that the fit might not be optimal. Here we compare the results and predictive performance using another common machine-learning method, the lasso.<sup>13</sup>

<sup>13</sup> The lasso is a regular OLS regression but with an added penalty that the sum of the coefficients cannot exceed a specified number. This helps to constrain the coefficients and reduce them to zero for the least informative variables, thus assuring that only the most important variables are included in the model. The penalty size is governed by a parameter,  $\lambda$ , which is often selected to minimize the out-of-sample error, or as one standard deviation less of the value that minimizes the out-of-sample error in an attempt to err to the side of parsimony and lower the risk of overfitting (Hastie et al., 2009).

**Table 2** Output from lasso regression

Variable	$\lambda_{opt}$ Coefficient	$\lambda_{optless1SD}$ Coefficient
Growth in GDP/capita	-0.1911	0.8506
*[datatype = income]	0.3540	--
*gini(0-100)	0.0162	--
*ln(population)(in millions)	0.0595	--
*headcountrate rate (0-100)	0.0004	--
*[East Asia and Pacific = Yes]	-0.0273	--
*[North America = Yes]	-0.0167	--
*[South Asia = Yes]	0.0200	--
*[Europe and Central Asia = Yes]	0.3114	--
*[Middle East North Africa = Yes]	-0.2834	--
*[Other high income = Yes]	-0.0928	--

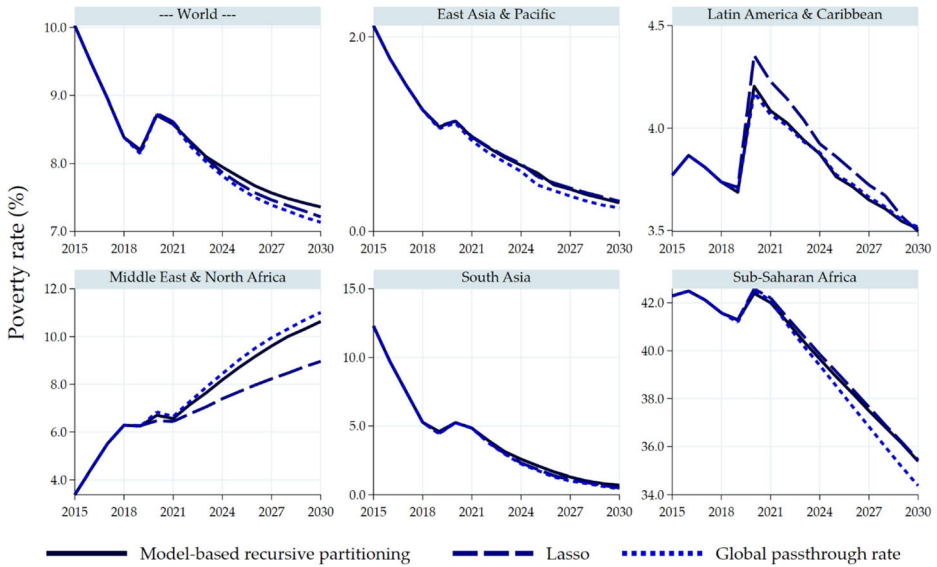
Note:  $\lambda_{opt}$  refers to the value of  $\lambda$  that minimizes the out-of-sample error from a 10-fold cross validation, while  $\lambda_{optless1SD}$  refers to one standard deviation less this optimal level. The last three rows refer to PovcalNet regions while the three other region rows refer to World Bank regions.

For the right-hand side of the lasso regression we use the same variables as before, but we convert the median, mean, GDP per capita, and population to log terms. This does not matter for model-based recursive partitioning since the splitting point for any of these variables would have been the same in log and non-log terms. Given that we are interested in the share of GDP per capita passed through to consumption, we include these variables as interactions with GDP per capita. Table 2 shows the results from the lasso, with  $\lambda$  selected to minimize the out-of-sample mean-squared error from a 10-fold cross validation, and one standard deviation less this optimal level.

When selecting the value of  $\lambda$  that minimizes the out-of-sample error, 10 variables matter for the passthrough rate. For example, with the optimal  $\lambda$ , the passthrough rate for a country that uses consumption, with a Gini of 40, a population of 10 million, a poverty rate of 20%, located in South Asia is predicted to have a passthrough rate of 0.62 ( $= -0.1911 + 0.0162 * 40 + 0.0595 * \ln(10) + 0.0004 * 20 + 0.0200$ ). When being slightly conservative and choosing the value of  $\lambda$  that is one standard deviation less than this, no variables matter for determining the passthrough rate, which then becomes 85% for all cases. We will refer to this as using a global passthrough rate below.

Compared with model-based recursive partitioning, both methods place a large emphasis on the datatype and the Gini index, and give Europe and Central Asia high passthrough rates. Model-based recursive partitioning makes use of the median, while the lasso uses the headcount, both of which are likely to contain much of the same information. Important differences lie in the emphasis on the population size in the lasso, and the large negative coefficient on Middle East & North Africa in the lasso, which is not (directly) picked up by the model-based recursive partitioning.

We can perform a cross validation to compare how well the three methods perform out of sample. Using the root mean square error, all three methods give an error between the predicted and actual annualized growth rates in mean consumption of 0.067 and the mean absolute deviation for all three methods is 0.040. The median absolute deviation is 0.025 for the global passthrough rate and model-based recursive partitioning but 0.024 for the lasso, suggesting a slight advantage to the lasso, but certainly not a difference that is statistically significant. The large difference between the



**Fig. 12** Distribution-neutral poverty projections with different passthrough rate methods. Note: Projected global and regional poverty rates measured at \$1.90 per person per day in 2011 PPPs under three different scenarios about how growth in GDP per capita is passed through to welfare observed in household surveys. The regional definitions follow the World Bank's classification, excluding the North America region and the Europe & Central Asia region where the share living below \$1.90 is minimal.

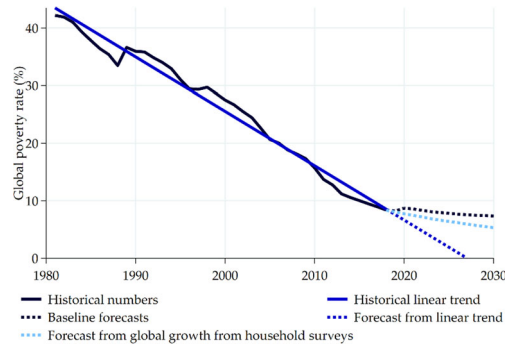
median absolute error and mean absolute error suggests that it is difficult to predict the passthrough rate for spells with extreme growth rates. In general, the out-of-sample errors combined with large differences between the two lasso models, as well as the fact that in the model-based recursive partitioning we were close to making no splits at all (the first split has a p-value of 0.041) suggests that the noise to signal ratio is very large.<sup>14</sup>

In general, the out-of-sample error could likely be improved with more sophisticated models, such as random forests, gradient boosting, or neural networks. Our aim here is to assess how much changes in inequality matter for poverty projections, and we think using more complicated models and testing them against each other would draw focus away from this main objective. We hope future research will compare additional methods that endogenize the determination of the passthrough rate.

The choice of passthrough rate method matters at the regional level but less so at the global level. Figure 12 shows the global and regional projections using our benchmark passthrough rate method as well as projections using the lasso and projections using a global passthrough rate (lasso with a value of  $\lambda$  that is one standard deviation less than the optimal value). Using model-based recursive partitioning gives slightly more pessimistic estimates at the global level with a forecasted global poverty rate of 7.4% in 2030 in contrast to 7.2% with the lasso model and the global passthrough rate.

<sup>14</sup> These results are consistent with Castaneda et al. (2019b) who use more sophisticated machine-learning methods and many more features to try to predict changes in mean consumption. They find that just using growth in GDP per capita performs nearly as well as selecting from hundreds of features.





**Fig. 13** Alternative projections to 2030. Note: Historical global poverty rates from PovcalNet (where the PovcalNet lining up method has been used for years where PovcalNet does not report a global poverty rate), and three projected global rates to 2030 assuming, (1) that the historical trend continues linearly, (2) that the historical growth in global mean consumption as observed in household surveys continues towards 2030, and (3) our baseline projections from earlier. Poverty line is \$1.90 per person per day in 2011 PPPs.

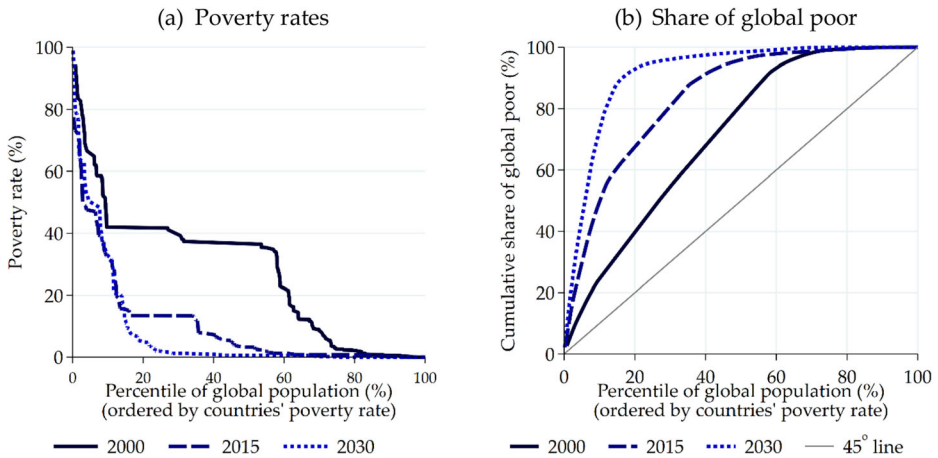
## 6 Discussion

Our projections suggest that getting close to the 2030 goal of ending extreme poverty will be unlikely. This contrasts with several of the scenarios in Ravallion (2013) that have global poverty below 3% in 2030. It is important to note that Ravallion (2013) was not a forecasting exercise, but primarily an attempt to set a goal for global poverty reduction that was neither too easy nor too hard, but could motivate effort.

When comparing our forecasts with the historical trend, the pessimism of our results might seem counterintuitive: Global poverty has decreased by almost 1 percentage point per year from 1981 to 2018, so how come we project a decrease in global poverty of about 1 percentage point *in total* from 2018 to 2030? If the historical trend continued linearly, one would find that global poverty would reach 0% well before 2030 (Fig. 13, based on Ravallion (2020)). Some of the divergence is due to COVID-19, yet even if the historical trend only were to continue from 2021, global poverty would still reach 0% by 2030. What explains this discrepancy?<sup>15</sup> One answer can be found by looking at another simple forecast: Using the annualized growth rate of the global mean from household surveys from 2000 to 2018 and applying this to the global distribution from 2018 onwards. Between 2000 and 2018, global mean consumption as observed in household surveys increased by 1.9% per year.<sup>16</sup> If we continue this towards 2030,

<sup>15</sup> One possible answer to this question which we will not consider here relates to incomparable estimates of poverty. Global poverty declined by around 4% as a result of China introducing imputed rent into its welfare aggregate (and switching to an integrated nationwide survey) (World Bank, 2016), and similar strong declines are expected when India goes from measuring consumption with a uniform reference period to a modified mixed reference period (World Bank, 2018). If changes to welfare aggregates systematically increase welfare, which the two examples above suggest and which would be the case if countries increase their capacity to include more elements in their consumption aggregate over time as they get wealthier, then our forecasts will overestimate poverty.

<sup>16</sup> The mean consumption observed in household surveys is taken from the lined-up global distributions in PovcalNet. When countries do not have surveys in 2000 and 2018, this lining-up relies on extrapolating distributions based on growth from national accounts, so in practice the growth rate we recover is based on a mix of growth observed in household surveys and growth observed in national accounts. This ensures that we look at a similar set of countries in both periods.



**Fig. 14** The world population ordered by country-level poverty rates. Note: The figure ranks the global population according to the poverty rate of their country of residence and plots this against the poverty rates and the cumulative share of global poor.

we get to a global poverty rate just above 5% – far from the linear projection and not too far from our baseline projection (Fig. 13).

This suggests that a non-negligible share of the global population remain quite far from the \$1.90 threshold. This is consistent with the finding from Ravallion (2016), that the consumption floor has risen little over the past decades, or equivalently, that there has been little progress in improving the welfare for the very poorest globally. Ravallion (2020) shows that historically it has been difficult for countries who have made remarkable progress in terms of eliminating poverty to go from a poverty rate of 3% to no poverty. All of this helps explain why reducing inequality rather than speeding up growth might be the most efficient way of getting these people to surpass the extreme poverty threshold.

Another (connected) reason for the deceleration of global poverty reduction is that the growth rate among the poorest countries has declined over the past decades (World Bank, 2018). This mostly boils down to growth in Sub-Saharan Africa being lower than growth was in China and India when the latter countries accounted for the bulk of the world's poor. With China and India making up a smaller share of the world's poor than in the past decades, progress needs to come from other places in order for the speed of poverty reduction to continue. Current growth forecasts do not suggest that this will happen.

In fact, when plotting country-level poverty rates in 2000 and 2015 (Fig. 14a), the largest changes come from countries ranked in the middle of the global distribution in terms of \$1.90 poverty. The population making up the poorest 10% of countries (ranked by their poverty rate) saw relatively small declines in headcount rates over this time period. Particularly, countries in long-term conflict and fragility have not managed to reduce poverty over the past decades (Corral et al., 2020). As our projections assume that this pattern continues, and as the scope for countries in the middle of the global distribution to further reduce poverty diminishes, this implies a leveling off in the speed of reduction of global poverty rates.

This also means that poverty is becoming more concentrated in a small number of countries (Fig. 14b).<sup>17</sup> In 2000, half of the world's poor could be found in countries

<sup>17</sup> We are thankful to R. Andrés Castañeda for making this point very clear to us.

making up 25% of the global population, while in 2015, half of the world's poor could be found in countries making up 10% of the world's population. If current trends continue, this number could fall to 5% by 2030, and countries making up 80% of the world's population will be almost free of extreme poverty while the remaining 20% will have poverty rates in high double digits.<sup>18</sup> The inability of 80% of the world's population to contribute significantly to global poverty reduction and the projected lack of progress in many of the countries in the bottom 20% explain why our forecasts suggest small gains over the next decade.

## 7 Conclusions

Using a global database covering 97.5% 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 Sustainable Development Goal of complete poverty eradication, will be difficult to reach by 2030. It also shows that these goals become more viable by reducing inequalities. 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, the shape of the distribution, and the growth incidence curve used. 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. In such cases, however, the poorest of the poor still receive a growth premium and thus the poverty gap and severity are reduced.

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 percentage point higher annual growth rates than expected. 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 richest individuals necessary to ensure that the bottom grows substantially faster than the mean is relatively small. In other words, the distributional changes simulated in this paper are technically feasible and highlight that pro-poor growth is crucial for reaching the poverty goals set by the global development community.

---

<sup>18</sup> As with all global poverty patterns, a lot of these patterns are driven by India and China, yet even when reproducing the figures without these two large countries, a qualitatively similar pattern emerges.

**Acknowledgements** The authors wish to thank R. Andrés Castañeda, Shaohua Chen, Francisco Ferreira, LaBlus Fah Jirasavetakul, Dean Jolliffe, Aart Kraay, Peter Lanjouw, Christian Meyer, Prem Sangraula, Umar Serajuddin, and Renos Vakis, as well as two anonymous referees and the editor for helpful comments and suggestions. 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. We gratefully acknowledge financial support from the UK government through the Data and Evidence for Tackling Extreme Poverty (DEEP) Research Programme, as well as the EFO No. 1340 (Measuring Poverty in a Changing World), and the Strategic Research Program (TF018888) for earlier versions of this paper.

## References

- Anand, S., Segal, P.: What do we know about global income inequality? *J. Econ. Lit.* **46**(1), 57–94 (2008)
- Atamanov, A., Jolliffe, D.M., Lakner, C., Prydz, E.B. Purchasing power parities used in global poverty measurement. World Bank Global Poverty Monitoring Technical Note, no. 5 (2018)
- Atamanov, A., Lakner, C., Mahler, D.G., Tetteh Baah, S., Yang, J.: The effect of new PPP estimates on global poverty: A first look. World Bank Global Poverty Monitoring Technical Note, no. 12 (2020)
- Birdsall, N., Lustig, N., Meyer, C.J.: The strugglers: the new poor in Latin America? *World Dev.* **60**, 132–146 (2014)
- Breiman, L., Friedman, J., Stone, C., Olshen, R.: Classification and regression trees. Taylor & Francis, Belmont (1984)
- Castaneda Aguilar, R.A., Lakner, C., Prydz, E.B., Soler Lopez, J.A., Wu, R., Zhao, Q.: Estimating global poverty in Stata: The PovalNet Command. World Bank Global Poverty Monitoring Technical Note, no. 9 (2019a)
- Castaneda Aguilar, R.A., Mahler, D.G., Newhouse, D.: Nowcasting global poverty. Paper presented at the Special IARIW-World Bank Conference. New Approaches to Defining and Measuring Poverty in a Growing World, Washington, DC, November 7–8 (2019b)
- Chandy, L., Ledlie, N., Penciakova, V.: The final countdown: Prospects for ending extreme poverty by 2030. Global Views Policy Paper 2013-04, The Brookings Institution, Washington DC (2013)
- Chen, S., Ravallion, M.: The developing world is poorer than we thought, but no less successful in the fight against poverty. *Q. J. Econ.* **125**(4), 1577–1625 (2010)
- Corral, P., Irwin, A., Krishnan, N., Mahler, D.G., Vishwanath, T.: Fragility and conflict: On the front lines of the fight against poverty. World Bank, Washington, DC (2020)
- Deaton, A.: Measuring poverty in a growing world (or measuring growth in a poor world). *Rev. Econ. Stat.* **87**(1), 1–19 (2005)
- Dhonde, S., Minoiu, C.: Global poverty estimates: a sensitivity analysis. *World Dev.* **44**, 1–13 (2013)
- Edward, P., Sumner, A.: Estimating the Scale and Geography of Global Poverty Now and in the Future: How Much Difference Do Method and Assumptions Make? *World Dev.* **58**, 67–82 (2014)
- Ferreira, F., Leite, P.: Policy options for meeting the Millennium Development Goals in Brazil: Can micro-simulations help? *Econ. J. Latin Am. Caribbean Econ. Assoc.* **3**(2), 235–280 (2003)
- Ferreira, F., Ravallion, M.: Poverty and inequality: the global context. In: Nolan, B., Salverda, W., Smeeding, T. (eds.) *Oxford Handbook of Economic Inequality*. Oxford University Press, Oxford (2009)
- Ferreira, F., Chen, S., Dabalen, A., Dikhanov, Y., Hamadeh, N., Jolliffe, D., Narayan, A., Prydz, E.B., Revenga, A., Sangraula, P., Serajuddin, U., Yoshida, N.: A global count of the extreme poor in 2012: Data issues, methodology and initial results. *J. Econ. Inequal.* **14**(2), 141–172 (2016)
- Foster, J., Greer, J., Thorbecke, E.: A class of decomposable poverty measures. *Econometrica.* **52**(3), 485–97 (1984)
- Hastie, T., Tibshirani, R., Friedman, J.: The elements of statistical learning: data mining, inference, and prediction. Springer Science & Business Media, Berlin (2009)
- Hellebrandt, T., Mauro, P.: The Future of worldwide income distribution. Working Paper Series 15 – 7, Peterson Institute for International Economics (2015)
- Higgins, M., Williamson, J.G. Explaining inequality the world round: Cohort size, kuznets curves, and openness. *Southeast Asian Stud.* **40**(3) (2002)
- Hillebrand, E.: The global distribution of income in 2050. *World Dev.* **36**(5), 727–40 (2008)
- Hoy, C., Samman, E.: What if growth had been as good for the poor as everyone else? Report (May). Overseas Development Institute, London (2015)
- International Monetary Fund: Fiscal policy and income inequality. IMF Policy Paper, Washington (2014)
- Jolliffe, D., Prydz, E.: Estimating international poverty lines from comparable national thresholds. *J. Econ. Inequal.* **14**(2), 185–198 (2016)

- Kakwani, N.: Poverty and economic growth with application to Côte d'Ivoire. *Rev. Income Wealth* **39**(2), 121–139 (1993)
- Karver, J., Kenny, C., Sumner, A.: MDGs 2.0: What goals, targets and timeframe? CGD Working Paper, Center for Global Development: Washington DC (2012)
- Kuznets, S.: Economic growth and income inequality. *Am. Econ. Rev.* **45**(1), 1–28 (1955)
- Laborde, D., Martin, W., Vos, R.: Poverty and food insecurity could grow dramatically as COVID-19 spreads. IFPRI Blog. Available at <https://www.ifpri.org/blog/poverty-and-food-insecurity-could-grow-dramatically-covid-19-spreads> (2020)
- Lakner, C., Mahler, D.G., Nguyen, M.C., Azevedo, J.P., Chen, S., Jolliffe, D.M., Prydz, E.B., Sangraula, P.: Consumer price indices used in global poverty measurement. World Bank Global Poverty Monitoring Technical Note, no. 4 (2018)
- Lakner, C., Milanovic, B.: Global income distribution: from the fall of the Berlin Wall to the great recession. *World Bank Econ. Rev.* **30**(2), 203–232 (2016)
- Lakner, C., Negre, M., Prydz, E.B.: Twinning the goals: how can promoting shared prosperity help to reduce global poverty? Policy Research Working Paper Series 7106, The World Bank (2014)
- Minoui, C., Reddy, S.: Kernel density estimation on grouped data: the case of poverty assessment. *J. Econ. Inequal.* **12**(2), 163–189 (2014)
- Ncube, M., Brixiova, Z., Bicaba, Z.: Can dreams come true? Eliminating extreme poverty in Africa by 2030. IZA Discussion Paper Series No. 8120 (2014)
- Pinkovskiy, M., Sala-i-Martin, X.: Lights, Camera... Income! Illuminating the national accounts-household surveys debate. *Q. J. Econ.* **131**(2), 579–631 (2016)
- Prydz, E.B., Jolliffe, D.M., Lakner, C., Mahler, D.G., Sangraula, P.: National accounts data used in global poverty measurement. World Bank Global Poverty Monitoring Technical Note, no. 8 (2019)
- Ravallion, M.: Growth, inequality and poverty: looking beyond averages. *World Dev.* **29**(11), 1803–1815 (2001)
- Ravallion, M.: Measuring aggregate welfare in developing countries: How well do national accounts and surveys agree? *Rev. Econ. Stat.* **85**(3), 645–652 (2003)
- Ravallion, M.: How long will it take to lift one billion people out of poverty? *World Bank Res. Obs.* **28**(2), 139–158 (2013)
- Ravallion, M.: Are the world's poorest being left behind? *J. Econ. Growth* **21**, 139–164 (2016)
- Ravallion, M. SDG1: The last 3%. Center for Global Development Working Paper No. 527 (2020)
- Ravallion, M., Chen, S.: Measuring pro-poor growth. *Econ. Lett.* **78**(1), 93–99 (2003)
- Rodrik, D.: The past, present, and future of economic growth. *Challenge* **57**(3), 5–39 (2014)
- Shorrocks, A., Wan, G.: Ungrouping income distributions. Working paper 2008/16, UNUWIDER (2008)
- Sumner, A., Hoy, C., Ortiz-Juarez, E.: Estimates of the impact of COVID-19 on global poverty. UNU-WIDER Working Paper 43 (2020)
- United Nations: Open working group proposal for sustainable development goals. United Nations, New York (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 (2018)
- World Bank: Comparability over time at the country level for international poverty measures. Data catalog, Washington DC (2019)
- World Bank. Global Shared Prosperity Database 2020. <https://datacatalog.worldbank.org/dataset/global-database-shared-prosperity>. Accessed June 1, 2020.
- Yoshida, N., Uematsu, H., Sobrado, C.E.: 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)
- Zeileis, A., Hothorn, T., Hornik, K.: Model-based recursive partitioning. *J. Comput. Graph. Stat.* **17**(2), 492–514 (2008)