Daniel Mahler (The World Bank), Christoph Lakner (The World Bank), Mario Negre (German Development Institute), Espen Prydz (The World Bank)

How Much Does Reducing Inequality Matter for Global Poverty?

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 of global extreme poverty (United Nations, 2014) or as reducing global extreme poverty to 3% of the world's population (World Bank, 2014). 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 called Model-based Recursive Partitioning 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 our set-up 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. 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.

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). 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.