universität wien

MASTERARBEIT / MASTER'S THESIS

Titel der Masterarbeit / Title of the Master's Thesis

„Will we eliminate poverty until 2030? An assessment based on the Growth Elasticity of Poverty"

verfasst von / submitted by

Fanny Helene Dellinger, Bakk.

angestrebter akademischer Grad / in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien, 2016 / Vienna 2016

Studienkennzahl lt. Studienblatt / degree programme code as it appears on the student record sheet:	A 066 913
Studienrichtung lt. Studienblatt / degree programme as it appears on the student record sheet:	Masterstudium Volkswirtschaftslehre / Masterprogramme Economics
Betreut von / Supervisor:	Univ.-Prof. Dipl.-Ing. Dr. Robert Kunst

Will we eliminate poverty until 2030? An assessment based on the Growth Elasticity of Poverty

Fanny Dellinger

April 2016

Abstract: I present a new and simple approach on how to forecast and replicate poverty trends based on the growth elasticity of poverty (GEP) and the inequality elasticity of poverty (IEP). I use an analytically derived elasticity and combine this with data from household surveys, i.e. mean income, inequality and growth rates are all taken from household survey data. The novelty of the approach is twofold: The income distribution is approximated by the Fisk distribution and instead of National Accounting System (NAS) data only household survey data is used. I perform an illustrative "test" of the model by replicating past poverty trends based on the survey data, and I find that the model generally performs well, (except for highly unequal middle income countries). I then apply the model to forecast future poverty trends and to assess whether the first target of the Sustainable Development Goals, i.e. to eliminate poverty until 2030, can be reached in the different world regions. My results show that in East Asia, the target will most likely be reached. Also in Latin America and South Asia, poverty is projected to be very low by 2030. It will be very difficult to reach the target in Sub-Saharan Africa.

Contents

List of Tables

List of Figures

1 Introduction

The last decades have brought extraordinary success in terms of reducing extreme poverty. According to UN-MDG Report (2015), in 1990 1.9 Billion people had to make a living from an income of less than 1.25 US$ a day. But by 2011, the last year for which comprehensive data is available, this number had dropped to 1.01 Billion and poverty is expected to drop even further, to 836 million in 2015. The level of income of 1.25 $ in 2005 PPP adjusted prices was chosen by the World Bank as the extreme poverty line. It is based on the median value of national poverty lines in the 10 poorest countries. The extreme poverty line was designated to measure poverty in poor countries. For middle or high income countries the line is far too low and relative notions of poverty, i.e. a poverty line expressed as a share of the median income, are of far greater importance. The extent of poverty reduction becomes much clearer if we focus on poverty rates. Those have declined by more than half, from 36% in 1990 to 15% in 2011, if we look at the world as a whole.

Will it be possible to keep up this pace of poverty reduction? If we managed to reduce the number of people with incomes below the extreme poverty line by more than half in just 15 years, will we be able to eliminate poverty by 2030?

The international community has set itself exactly this goal. The first target of the Sustainable Development Goals (SDGs), which were ratified by the UN General Assembly in September 2015, is to eliminate extreme poverty by 2030.

But assuming that poverty will continue to fall at this exact pace, i.e. linearly extrapolating a poverty trend, is far from perfect. While the poverty decline we have witnessed so far might seem linear, in fact, it is not. The shape of the income distribution has to be taken into account especially when it comes to reducing poverty to zero. This automatically implies a slowing down of the pace of poverty reduction when the poverty rate becomes low (Chandy et al., 2013; Ravallion, 2013).

Poverty reductions are overwhelmingly associated with income growth (Chen and Ravallion, 2001; Adams, 2004). Assessing poverty trends then becomes a question of quantifying the relation between income growth and poverty. A popular tool for doing so is the growth elasticity of poverty (GEP). It shows by how much percent poverty declines if average incomes grow. It is thus a simple but nonlinear tool for assessing poverty trends. Ideally it provides us with a means to make informed statements about future poverty reduction. The GEP is especially useful for cross country comparisons as it does not require working with the full distributions (Bourguignon, 2004). In particular, we can answer the question of whether current growth rates are sufficient to achieve the goal of eliminating poverty by 2030? Can we achieve this goal primarily relying on growth, or does it require policies of (targeted) redistribution? Of course, the GEP can't answer all questions, but it helps us to make informed statements.

In my paper, I present a new simple approach on how to forecast and replicate poverty trends based on the growth elasticity of poverty (GEP) and the inequality elasticity of poverty (IEP). The two innovations I introduce are first, a switch from the Lognormal distribution to the Fisk distribution for approximating the income distribution when analytically deriving the GEP. So far, economists have mainly relied on the Lognormal distribution for approximating income distributions. This practice has been criticised as the overly reliance on just one distribution function could introduce a bias in empirical work. Bresson (2009) has demonstrated that especially in middle income and unequal countries there is a danger of upward bias of the GEP.

Second, I base my forecasting calculations on household survey data and not on data from the National Accounting System (NAS), e.g. GDP or private final consumption expenditure. Many authors such as Deaton (2005), Ravallion (2003a) or Sundaram and Tendulkar (2003) have convincingly argued and by now it is established practice at the World Bank that poverty data should be assembled based on household surveys. Household surveys and NAS aggregates tell quite different stories about the living situation of the people in developing countries. There are important discrepancies in terms of levels and growth rates of income and private consumption expenditure. Thus, it is of uttermost importance that the correct data is used for forecasting poverty trends.

The structure of my paper is as follows: In section 2, I review the literature on the growth elasticity of poverty highlighting the different approaches of how it can be derived. Section 3 analyzes questions related to the measurement of poverty and I draw conclusions from this literature on how to better interpret the different results of the GEP literature. In section 4, a simple model of poverty reduction, the derivation of the GEP-formula based on the Lorenz curve, and the specifics of the Fisk elasticity are presented. In section 5, I review data issues related to the discrepancies between NAS data and household survey data. Section 6 on results has two parts. First, I informally "test" my approach by replicating past poverty trends. I find that the approach yields good trend approximations for poor countries. Middle income countries with large shifts of inequality are more difficult to handle. This clearly shows the limits of the approach. Second, I produce some tentative poverty trend projections for four world regions. I focus on the regions with the largest number of poor people in the world and I try to assess whether the Sustainable Development goals will be achievable. Section 7 compares my results to other related studies while section 8 concludes.

My work touches upon the bigger question of whether we are actually able to measure economic development, inequality and societal well being. This paper highlights some of the caveats of conventional economic measures like GDP, private consumption expenditure or the Gini coefficient. At least, we need to be aware of these.

2 Growth Elasticity of Poverty

The growth elasticity of poverty (GEP) is a useful tool for quantifying the relationship between income growth and poverty. The elasticity tells us in percent by how much poverty will change, if income increases by 1 percent. As we expect poverty to fall, if incomes increase, the elasticity will in general be negative. Also, a higher absolute value of the elasticity means that the effect of income growth on poverty is stronger. Formally we can write:

$$GEP \equiv \eta_{Pov}^{Inc} = \frac{\Delta Pov/Pov}{\Delta Inc/Inc}. \tag{1}$$

At the end of the 1990s and beginning of the 2000s there was a lively academic debate with many researchers trying to provide sound and comprehensive estimates of the GEP. Usually, researchers tried to quantify the relation between income growth and changes (hopefully declines) of the poverty headcount, a poverty measure, which gives the proportion of people below a poverty line. Reported values varied quite a lot. See Adams (2004) for a review. Maybe prematurely, in 2000 the World Bank concluded that on average the GEP is -2 in its study “Attacking Poverty” (World Bank, 2000). For every percentage point that GDP increases, the poverty headcount will decrease by 2 percent.

2.1 A review of the literature on the GEP

The year 2000 brought the Millennium Declaration of the UN and the Millennium Development Goals (MDGs). Target 1 of the MDGs called for a halving of extreme poverty in the world. This was defined to mean a decline by more than half of the percentage of the world's population below the extreme poverty line. This line was set at 1 US$ in the year 1985 and later updated to 1.08 US$ in 1993 purchasing power parity (PPP) adjusted US-Dollars. For more on poverty lines and the rationale behind this specific poverty line, see Chen and Ravallion (2001). A large revision of past poverty numbers took place in 2008, when the poverty line was adjusted to 1.25 US$ and new PPP adjustment factors anchored to the year 2005 came into play (Chen and Ravallion, 2008). Today, the extreme poverty line is set at 1.9 in 2011 PPP US$. The proclamation of the MDGs and in particular target 1 sparked the interest of researchers into forecasting poverty trends. Also, throughout the 1990s household survey data (of nationally representative surveys) became available for more and more countries. And, the data seemed to tell a good story: Poverty was declining. But to what extent and at what speed?

The growth elasticity of poverty is a measure which helps to answer this new question. Based on a sound estimate of the GEP it is possible to assess current and forecast future trends of poverty and to assess whether a poverty reduction goal can be met in time, as was demonstrated by Besley and Burgess (2004).

It was difficult to arrive at a consensus on the size of the GEP. The early estimates covered a wide range, from below 1 (Besley and Burgess, 2004) to about 3 (Ravallion and Chen, 1997).

Different approaches of how to derive the GEP have been suggested in the literature. In the 1990s and late 2000s many studies tried to estimate the relationship by common econometric tools (Bourguignon, 2004). Researchers often relied on data sets containing growth spells, which show the levels of income and of some inequality measure at all points in time when poverty estimates are available. The methodology was largely developed by Ravallion and Chen (1997).

Some confusion in the literature over diverging results stems from the mixing of different data sources. The data on income (or consumption) can either be taken from household surveys or from the national accounting system (NAS). The former are large nationally representative household surveys of either income or consumption, which serve the purpose of gathering both distributional and poverty data. The latter is used for deriving key macro-economic figures like GDP. Ravallion and Chen provide a theoretical argument for why the usage of the survey means, (i.e. the average income or consumption expenditure of those surveyed) might actually be preferable (Ravallion and Chen, 1997). Household survey design varies from country to country, (e.g. either income or consumption is surveyed, different recall periods are used and the listings of consumption items, which are asked for, differ) but these surveys are the basis for the poverty figures. This is bound to introduce some error in a cross country comparison of poverty figures. But when it comes to investigating the relationship between income growth and poverty in a panel, fixed effects framework, this might not be so problematic. The survey means will also suffer from survey design related comparability problems. As in any model with an explanatory variable which is measured with error, this will introduce an attenuation bias and it will be harder to find any significant relationship. Now, it is quite safe to assume that, on average, income growth lowers poverty. If average incomes in the survey are overestimated, poverty will be underestimated and vice versa. This introduces another bias, a common survey bias, which actually makes it more likely to find a significant relationship. Thus, we have two different biases with opposite signs. If we are willing to make further assumptions and to impose some structure on the errors, these biases are exactly offsetting.

While there are econometric reasons for preferring survey means, arguably we might be more interested in how poverty changes with respect to GDP growth. (Also, the survey means themselves are usually not publicly available.) Besley and Burgess (2004) estimate the GEP in a fashion not unlike the one proposed by Ravallion and Chen, but using GDP instead of survey means. In absolute terms, their results are much lower (close to and for poor regions clearly below 1) than the ones of Ravallion and Chen (which are on average -3 when the poverty line is set at 1 US$ per day). The authors do not discuss the reasons for this discrepancy.

Adams (2004) also presents estimates of the GEP. He explicitly discusses the difficulty related to different data sources on income. Adams does not reach a clear cut conclusion of how the usage of either survey means or NAS growth impacts the GEP. But this might be due to his heterogeneous data set, often blurred by maybe unconventional episodes stemming from the transition economies in Eastern Europe and Central Asia.

The econometric approach has come under attack by Bourguignon (2004), who argues that a linear or close to linear specification of the relationship between growth and poverty is not adequate. Instead, he stated that growth and changes in the poverty headcount are linked via a highly non-linear identity relationship. This relationship is characterized by inequality but also by the distance between the poverty line and the mean of the income distribution. We can expect countries with higher mean income to also display a GEP, which is higher in absolute terms. More unequal countries will, in general, have lower GEPs. If we are willing to make some assumptions on the form of the income distribution, i.e. that it can be approximated by a known cumulative distribution function $F(x)$, the question of deriving the GEP becomes a purely mathematical one. The poverty headcount can then be given by the expression $F(z)$, when z is the poverty line. It is possible to exactly calculate how $F(z)$ changes, if all incomes grow at the same rate, but the poverty line does not change. Bourguignon (2004) proposes to approximate the income distribution by the Lognormal distribution function. Then the GEP is encapsulated in the formula (2). H_t stands for the poverty headcount in the year t and μ_t stands for mean income in year t.

$$GEP = \eta^{\mu}_{H,t} = \frac{\Delta H_t}{\Delta \mu_t} \frac{\mu_t}{H_t} \equiv -\frac{1}{\sigma_t} \times \frac{\varphi_t(\frac{log(\frac{z}{\mu_t})}{\sigma_t} + \frac{1}{2}\sigma_t)}{\Phi_t(\frac{log(\frac{z}{\mu_t})}{\sigma_t} + \frac{1}{2}\sigma_t)}. \tag{2}$$

The GEP is a function of the log standard deviation σ_t, the mean income μ_t and the poverty line z. In absolute terms, the elasticity is decreasing with inequality σ_t and increasing with μ_t. Also, Bourguignon showed how to calculate the effect of changes in the inequality parameter of the distribution function, and presented a formula for the elasticity of poverty with respect to the log standard deviation of income. Bourguignon's results have become quite popular, as they require very little data input and allow to, for instance, decompose observed poverty reductions into the change due to income growth and the change due to changes in inequality (Kalwij and Verschoor, 2005). More recently, the approach has come under critique.

Bresson (2009) argued that the reliance on the Lognormal distribution for deriving the analytical expression of the elasticity is not optimal. Indeed, it is likely that this assumption introduces a bias in the derived elasticities, and especially for middle income countries which are not very unequal, (to be precise, when the Gini index is lower than 42% to 45%), these elasticities seem to be too high. At the same time Bresson does not question the existence of an identity relationship linking the income level and inequality to the growth elasticity of poverty, but he suggests working with more sophisticated distributional forms, (for instance distributions which rely on more than just two parameters), if one is truly interested in deriving reliable elasticities.

The concept of GEP was challenged by Ram (2006). Between 1990 and 2002, GDP in the developing world grew by 25%, so if the growth elasticity of poverty had indeed been of the order of -2, the decline of poverty should have been much stronger. The values often put forward in the literature cannot be reconciled with actual poverty and GDP growth trends. Developing this idea further, Lenagala and Ram (2010) analyse a new dataset on poverty by calculating past trends of poverty and growth for different regions, and they define a "direct" measure for the GEP, namely the annual percentage change of poverty divided by the growth rate. They find relatively low elasticities. In particular, for higher poverty lines, e.g. the 2 US$ or the 2.50 US$ per day poverty line, the elasticities are strikingly low, especially in India.

2.2 A comparison of the three approaches

Which approach is most adequate for deriving the GEP? Dellinger and Diprossimo (2015) ran a comparison of the three approaches outlined above. Their objective was to highlight how the choice of any particular approach influences the outcomes. They replicated all three methods and calculated the GEP for the different world regions, using the poverty data provided by the MDG data base. The data covers a time period from 1990 to 2012. Unlike most studies cited earlier, the 2000s are also covered. Another innovation with regard to earlier work is the weighing of country observations by population size. For the analytical approach, for each country and each year the elasticity is calculated according to formula (2) and then a population weighted average of the different regions is formed. For the direct elasticity, second order polynomial trends of poverty and GDP are calculated for the different regions, thereby taking into account country size. Then for each year, the ratio of the growth rates of the two trends is formed. The mean and the median of these ratios is calculated. Table 1 summarizes the results of their work.

Several conclusions can be drawn. First of all, the analytical elasticity usually gives the highest result. (The only exception is Latin America, where the mean of the direct elasticity is highest.) Also, and as expected from the

	East Asia	South Asia	Latin America	Europe and Cent. Asia	N. Africa and M. East	Sub-Saharan Africa
Regression based	-1.43	-0.71	-1.87	-3.31	-1.9	-0.47
Analytical elasticity	-2.21	-1.55	-2.16	-5.63	-3.12	-0.78
Direct elasticity Mean	-1.62	-1.04	-2.18	-4.84	-1.07	-0.35
Direct elasticity Median	-1.22	-0.68	-2.12	-2.12	-0.98	-0.43

Table 1: The growth elasticity of poverty. Observations weighted by country size

previous discussion, we find that regions where middle income countries prevail, like Europe and Central Asia, North Africa and the Middle East and Latin America display higher values of the GEP. The only exception is the analytical elasticity which yields a higher value for East Asia than for Latin America. This is probably due to the high Gini coefficients of Latin America. The lowest values of the GEP are consistently found in Sub-Saharan Africa, where incomes are very low and inequality is high. It helps to emphasize that a high GEP in middle income countries does not mean that poverty reduction will be faster there. In fact, this slightly higher elasticity only partly offsets the decrease in the pace of poverty reduction, which is usually observed when countries become richer. Let us consider a simple example. If the poverty rate is 50% in a poor country, and 5% in a middle income country, and the GEP is -1 in the poor country, and -5 in the other, then income growth of 2% will lower the poverty rate of the poor country by 1 percentage point. The poverty rate of the middle income country will fall by 0.5 percentage points, although the absolute value of the GEP is much larger there. If the middle income country GEP were also -1, poverty would fall by just 0.1 percentage points.

The previous discussion has shown that it is to be expected that the analytical elasticity is slightly higher in absolute terms than the other two. There is some basis for believing that it might be overestimated. First, the criticism of Bresson (2009) has to be taken into account. Second, there are problems of data comparability. For poverty measurement, the World Bank exclusively relies on the data from household surveys, which show far lower levels and lower growth rates than the data from the National Accounting

System (NAS). The pros and cons of measuring poverty by relying on NAS data have been extensively discussed in the literature (Deaton, 2005; Ravallion, 2003a; Dhongde and Minoiu, 2011). These arguments might well also apply when it comes to deriving the GEP. If so, simply using GDP data for deriving the analytical elasticity is not adequate.

Finally, the results of Dellinger and Diprossimo (2015) show quite clearly, that the GEP increases with the mean level of income. If we want to forecast development trends which hopefully imply growing incomes and declining poverty, assuming the GEP to be stable is almost certainly a mistake. The analytical elasticity enables researchers to forecast development trends by taking into account how growing incomes impact the growth-poverty relationship. For assessing and forecasting progress towards the Sustainable Development Goals an analytically derived elasticity seems most promising. But first the two problems of the analytical elasticity outlined above require some in depth analysis.

3 The tricky business of measuring poverty

How did the poor fare in the last two to three decades? Did they actually profit from globalization? Or have they been left behind? In order to give an answer to these questions, it is necessary to measure and quantify poverty and poverty trends. Sound knowledge of the GEP can be useful here, but of course the debate goes much further. Also, it will become clear that the main discrepancy between the different estimates of the GEP, namely that the regression based and the direct approach yield lower estimates than the analytical approach, is directly linked to a recurrent theme of poverty measurement: the difference between data from the national accounting system and household surveys.

As for the poor and globalization - while I am unable to discuss this topic in detail, I would like to refer to the interesting overview of the topic by Ravallion (2003a). He argues that while undeniably the poor profit from growth and are hit hard by recessions, it is too easy to claim that the poor unequivocally profit from globalization and that distributional effects can be ignored.

3.1 The contentious figures on poverty

Unfortunately, it is all but straightforward to measure poverty. The historical practice and many of the most relevant difficulties have been explored in detail by Deaton (2006).

The World Bank has an important role to play, as the organization is in charge of the official world poverty count. It provides internationally comparable data on the proportion of the population below specific poverty lines. The extreme poverty line is set to reflect poverty in the poorest countries (Deaton, 2006; Chen and Ravallion, 2001).

These official figures attract a lot of criticism from all sides of the political spectrum (Ravallion, 2003a). An extreme position is taken by Bhalla in his book "Imagine There's No Country: Poverty, Inequality and Growth in the Era of Globalization" (Bhalla, 2002). Therein the author argues that globalization has been overwhelmingly beneficial for the poor. He presents his own figures on poverty based on more than 900 income distributions with some dating back to the 1950s. Mean income is not taken from these distributions, instead conventional GDP, uniformly deflated by 15% across all countries, is used. The numbers on poverty are then generated using the mathematical properties of poverty measures. These numbers contradict the World Bank figures and show a much larger decline of poverty. According to Bhalla, already in the year 2000, when the goals were set, the millennium development goals would have been met.

Bhalla's work has been heavily criticised by the World Bank economist Ravallion (2002) for not applying any quality standards for the data, often not adjusting for PPP and inflation correctly, making incredible assumptions and general technical insufficiency.

Sala-i Martin (2006) also calculated figures on absolute poverty in the world, based on the world income distribution which he assembled. Sala-i Martin (2006) was more careful when it came to applying quality standards to country distributional data. He also came to the conclusion, that the poverty figures of the World Bank are too high.

Where does this discrepancy come from? Most importantly the figures of Bhalla (2002) and Sala-i Martin (2006) are based on a combination of inequality measures and NAS data, while the World Bank figures exclusively rely on nationally representative household surveys of either consumption expenditure or income. The levels of private household consumption from the surveys is generally lower than from the NAS accounts (Deaton, 2005; Ravallion, 2003b). This issue is crucial for poverty measurement. One of the reasons for why these measures diverge is that well-situated households are less likely to participate in household surveys and if they do, there are some indications that they might not disclose all of their consumption or income.

Bhalla and other authors such as Sala-i Martin (2006) construct income distributions for each country (for which some data is available) and then average over these country distributions to arrive at a world distribution of income, which enables them to subsequently derive poverty counts. The mean of the country distribution is given by country GDP (taken from the Penn world tables in the case of Sala-i Martin (2006)). For within country inequality distributional data is used, which is taken from household surveys.

Ravallion (2003a) provides a simple and convincing explanation, of why this is a bad idea. His main point is that while it is perfectly reasonable to assume that household survey means are too low, there is no reason whatsoever to assume that at the same time, the surveys give the correct inequality measure. Instead, it is far more likely that the surveys underestimate both the level of average household consumption and the extent of inequality within a country. Indeed, a study from Latin America seems to confirm this. Analyzing a large set of Latin American household surveys, Székely and Hilgert (1999) show that the 10 highest individual incomes measured per survey are of the order of magnitude of typical manager incomes of medium to large firms and thus far lower than what we would expect from a truly representative sample (Székely and Hilgert, 1999).

I would like to quote an illustrative example given by Ravallion (2003a) literally:

"To see why anchoring poverty measures to the national accounts can go so wrong, consider the following simple example. The true but unobserved distribution of income is (say) 1,2,3 (person 1 has an income of 1, person 2 has income 2, person 3 has 3). The poverty line is slightly above 1, so the true poverty rate is 1/3. We do a survey, and the three people respond that their incomes are 1, 1.5 and 2. This also gives the right poverty rate. However, the survey underestimates the true mean; the survey mean is 1.5. Now let's assume (for the sake of argument) that the national accounts do give the right mean of 2. If we assume that the survey under-estimation is distribution-neutral then we multiply all three incomes by 4/3. The "corrected" incomes are 1.3, 2 and 2.7 - implying that there is no poverty. We get the mean right, but the poverty measure is way off the mark."

While this example is highly simplistic, Ravallion further argues that it might nevertheless capture the core of the problem quite well. In a study for the US Mistiaen and Ravallion (2003) analyze whether unit non-response is linked to income. Based on data of survey compliance and income at the district level, they find that there is a bias introduced in the surveys due to selective unit non-response by the rich. It is also possible to correct for this bias. While little correction is necessary at the lower end of the distribution, consumption should be scaled up by 30% to 50% for the top quintile. In a study for India Sundaram and Tendulkar (2003) report that for items making up approximately 75% of the consumption expenditure of the poor, the difference between NAS private households consumption expenditure and consumption measured in the surveys is actually very small. Both studies indicate that the discrepancy between survey and NAS data is more of a problem when measuring the consumption of the rich. The surveys seem to work reasonably well when it comes to determining the consumption expenditure of the poor. Also, both studies imply that there might be problems with measuring inequality.

Much the same debate has taken place a few years later regarding Africa. Sala-i Martin and Pinkovskiy (2010) claimed that poverty in Africa was falling much faster than official numbers indicate and that inequality was actually improving. Based on a similar methodology as Sala-i Martin (2006) they construct an African income distribution based on GDP data from the Penn World Tables and Gini coefficients from the WIID - World Income Inequality Database. This allows them to calculate poverty rates over time. Sala-i Martin and Pinkovskiy (2010) state that target 1 of the MDGs, halving the poverty rate from 1990 to 2015, will be achieved just two years late in Sub-Saharan Africa. Their conclusions have been harshly criticised by McKay (2013). First, McKay pointed out that the data on inequality in the WIID is not always consistent and insufficient for making any strong claims regarding the trend of inequality in Sub-Saharan Africa. Also, McKay voiced scepticism regarding the choice of using GDP as the mean of the African income distributions. This is especially hard to defend taking into account that the quality of household surveys has greatly improved in the last ten years or so, while NAS figures are still often of dubious quality (Jerven, 2009; McKay, 2013).

The widely diverging claims about the extent of poverty in the world have also been analysed by Dhongde and Minoiu (2011). They emphasize the importance of choosing either the NAS aggregates or the survey means for the calculations. Conducting a sensitivity analysis they quantify the effect of this choice on global poverty figures by calculating global poverty in 1995 and 2005 based either on survey means, NAS consumption or NAS GDP. Poverty is given by a function of the Lorenz curve, which is kept constant, but whose mean is anchored to one of the three options outlined above. For the 1 US$ a day poverty line, they find that the global poverty rate in 1995 was 29%, based on the survey means, 5.9% based on NAS consumption and 1.4% based on GDP. By 2005 poverty had declined across all three means. The survey means gave a poverty rate of 24.3%, NAS consumption yielded poverty of 1.7% and GDP gave a poverty rate of 0.9%. Obviously, the choice of the mean drives the results. And it is little surprising that authors who calculate poverty based on GDP or NAS consumption report much lower poverty figures.

3.2 Why do NAS aggregates and household survey data differ?

There is a sizable and well documented divergence of survey means and NAS data (Deaton, 2005; Ravallion, 2003a; Bhalla, 2003; Dhongde and Minoiu, 2011). Especially in India, where, approximately, one third of the world's poor live, the difference between surveys and national accounts is startling. The ratio of consumption as measured in the surveys to consumption taken from the national accounts was just 0.56 in the year 2000 (Deaton, 2005). (It should be emphasized that the Indian household surveys are usually considered to be of high quality (Deaton and Kozel, 2005).) In addition in most countries the divergence between surveys and NAS is further widening, suggesting different growth rates of NAS aggregates and household consumption as measured by surveys (Dhongde and Minoiu, 2011).

Why do NAS aggregates and survey means systematically differ? Some discrepancies between surveys and NAS aggregates are to be expected due to issues of definition and coverage (Sundaram and Tendulkar, 2003). The closest NAS item to what is measured in a household consumption survey is private final consumption expenditure. This is calculated as a residual of the commodity flow method. First an overall estimate of production is made and adjustment for foreign trade takes place; then government consumption and absorption by firms (i.e. changes of inventories) are subtracted. What is left is essentially NAS consumption. The consumption of not-for-profit non-governmental organizations (including among others religious institutions, political parties, schools run by NGOs etc.) is counted as household consumption and it cannot be disentangled from the latter. (For India, Sundaram and Tendulkar (2003) point out that the NGO sector in India is growing quickly, and thus contributing to the diverging growth rates.) At the same time, household surveys exclude people in institutional housing (i.e. orphanages, prisons, hospitals...) and the homeless. Finally, two important and fast growing consumption items are included in the NAS calculation, which do not feature in the surveys: imputed rents from owner-occupier dwellings and the imputed value of financial services intermediation. Estimating the former for the poor is not straightforward, especially in (rural) areas where there is no rental market. The latter is calculated as the difference between fees charged by banks and insurance companies and the costs of the industry, as it is assumed that people derive an extra benefit from the availability of financial services.

All of these differences contribute to the divergence of NAS aggregates and survey means and, so far, we did not even consider (measurement) errors in the data. Unfortunately, both the surveys and the NAS aggregates suffer from errors.

As pointed out before, with the surveys unit non-response especially by the rich is a problem, which leads to downward biased estimates of mean consumption or income. If inequality is rising and the rich are not properly sampled, this might also account for the diverging growth rates of NAS aggregates and survey means (Deaton, 2005).

Also survey design matters. For instance, there is no unified approach to issues such as who to ask in a household, the level of disaggregation, etc. Recall periods are of particular concern. Should people be asked to report their weekly consumption or their monthly consumption of non-durable goods? Shorter recall periods typically lead to more consumption being reported (Deaton and Kozel, 2005). In India, the surveys have been operating with 30 days recall periods, while in many other countries 7 day recall periods are common. Weekly recall periods give a consumption estimate which is about 23% higher than the one for monthly recall periods. For a wide range of goods they are also more accurate. But for a few important commodities, especially rice, which account for a large share of the consumption of the poor, it turned out that the monthly recall period is actually better suited (NSSO Expert Group, 2003; Deaton and Kozel, 2005).

At the same time there are serious problems with NAS aggregates as well. First of all, as emphasized by Deaton (2005), they were designed to track money, not people and especially not the poor. As consumption in the NAS is calculated as a residual, previous errors (of production estimates, changes in inventories, intra-industry consumption and investment, intermediate products...) all accumulate in the NAS consumption figures. As Sundaram and Tendulkar (2003) point out, there is no obvious reason why these errors should cancel out. Especially intermediate production is problematic. Its calculation relies on industry surveys, ratios and approximations which are often not much more than educated guesses of the statisticians involved in producing these numbers (Sundaram and Tendulkar, 2003; Deaton, 2005; Minhas, 1988). The industry surveys are conducted on an irregular basis and they are sometimes decades old. Every few years, when more information becomes available, large revisions of NAS aggregates are necessary (Sundaram and Tendulkar, 2003). The paucity of these ratios has already been pointed out by Minhas (1988) for India. Unfortunately, since then, the situation has not changed much, according to Sundaram and Tendulkar (2003); Deaton and Kozel (2005). Back in the 1980s, criticism of the NAS aggregates led to the decision to not adjust the means of the household surveys to fit the level of NAS consumption for measuring poverty in India (Deaton and Kozel, 2005).

Another source of a possible upward bias is found in the transition of a subsistence economy to a market economy. Home production is not captured in GDP. But in this process of transformation home production is replaced by market production, for instance, because people buy more food outside and cook less at home, and this will increase GDP without necessarily increasing consumption (Deaton, 2005). Many such processes characterize the transformation the developing world is undergoing, and this is a source of bias for their GDP growth.

NAS consumption is especially problematic, but even GDP is not as reliable, as we are used to think, in particular in the developing world. Jerven (2009) assessed the reliability of GDP data in Africa. He compared GDP rankings of Sub-Saharan African countries when GDP data is taken from three different publicly available data bases, i.e. World Development Indicators, Penn World Tables and Maddison. The different data bases give very different GDP data for the Sub-Saharan African economies, which also affects the relative ranking of the economies. These are found to differ greatly and they are not stable over time. Of course, this is bad news for everybody engaging in econometric work about African countries.

To sum up, most economists accept the World Bank's methodology of relying exclusively on survey data for constructing poverty estimates. First, it serves to reiterate the argument of Ravallion (2003a) that although there is a good chance that the surveys show means with a downward bias, at the same time it is almost certain that they also underestimate inequality. Second, NAS data is far less reliable than what we are usually hoping for. This is especially true for developing countries where the administrative capacity to collect the necessary data might be lacking, but problems with NAS accounts also affect the developed world. Stiglitz et al. (2009) have argued convincingly for caution when using GDP. It is not a universal welfare measure and policy makers should not focus exclusively on increasing GDP. There are many examples of what may go wrong when relying on GDP only, and why comparisons of countries based on GDP alone might be misleading. Goods and services provided by the government to the households are included in GDP, but there is no market value i.e. price, at which these services can be counted. Instead, their value is determined by the input costs. So, if a government improves efficiency and provides good services at lower costs, this decreases GDP. The US, for instance, has a very expensive health system but health outcomes, e.g. in terms of life expectancy, are below the OECD average. If the health sector in the US and France were given the same weight in GDP, the difference in GDP per capita between the US and France would be reduced by one third (Stiglitz et al., 2009). Such aspects of GDP might matter a lot for cross-country comparisons and we should not be oblivious to them.

3.3 Implications for the GEP debate

Some very important conclusions can be drawn from the debate about measuring poverty for the study of the GEP. Most importantly, if NAS data and household survey data differ substantially, we need to clarify what data sources are used for deriving the GEP.

Different claims made about the size of the GEP become easy to interpret. If GDP grows at a faster rate than mean income from household surveys, then a regression explaining poverty changes by growth will yield an (in absolute terms) lower coefficient when using GDP growth than when using survey income growth as explanatory variable. So, the elasticity of -3 found by Ravallion and Chen (1997) and the elasticity of -0.7 found by Besley and Burgess (2004) also have to be seen in this light. The former authors used the growth of survey means as explanatory variable, the latter the growth of GDP.

Also, the differences with regard to the analytical elasticity become clearer. The analytical elasticity needs mean income and an inequality measure as inputs. GDP per capita and, the more appropriate comparison value for most surveys, NAS private consumption expenditure show far higher levels than mean income and consumption from the household surveys. If for calculating the GEP we use NAS data, we will get higher values of the GEP than if we had used the survey data. In case of the analytical elasticity, just relying on NAS data can be dangerous - we combine levels with a potential upward bias with inequality data with a likely downward bias. Any analysis based on this combination of data, e.g. Kalwij and Verschoor (2005), has to be viewed with great caution.

Finally, also the criticism of Lenagala and Ram (2010) and Ram (2006) can be put into a new perspective. They compare the value of -2 of the GEP presented by the World Bank with actual trends of GDP and poverty. But the World Bank's estimate of the GEP is mainly based on the work of Ravallion and Chen (1997) and others who derive the GEP as the relationship between poverty and household survey growth of incomes or consumption, not between poverty and GDP growth. Again, lower growth of household survey means implies a higher GEP. Their criticism should not simply be dismissed. Their results can also be interpreted in another way - not the value of the GEP is problematic, instead the fact that the incomes of the poor grow at a much lower rate than shown by GDP growth combined with conventional inequality measures is the real reason for concern.

What could a new and comprehensive estimate of the GEP look like? An analytically derived elasticity has some clear advantages when it comes to forecasting trends. But of course, the objective should be to derive results which are in line with the official poverty counts. Thus, it becomes essential to compute the GEP exclusively from household survey data. And finally, any forecast of poverty trends should be based on trends in growth of household survey means.

4 A simple model of poverty reduction

I set up a simple model of the annual change of the poverty head count based on the GEP and the inequality elasticity of poverty (IEP). Required input data are the poverty headcount in the year t, H_t, mean income, μ_t, and the its growth rate, the Gini coefficient, $Gini_t$, and its rate of change. I define the GEP as $\eta^{\mu}_{H,t}$ and the IEP as $\eta^{Gini}_{H,t}$.

$$H_{t+1} = H_t * (1 + \frac{\Delta\mu_{t+1}}{\mu_t} * \eta^{\mu}_{H,t}) * (1 + \frac{\Delta Gini_{t+1}}{Gini_t} * \eta^{Gini}_{H,t}). \qquad (3)$$

Besides the elasticities, the growth rate of μ, $\frac{\Delta\mu_{t+1}}{\mu_t}$, and the rate of change of the Gini coefficient, $\frac{\Delta Gini_{t+1}}{Gini_t}$, are essential for determining the new poverty level. First, I will analyze the elasticities in detail, and then I will turn to μ_t. Properly defining μ_t is decisive for making the simple model presented in equation (3) work.

4.1 A general formula for GEP and IEP

It is possible and fairly straightforward to derive a general formula for the GEP from the Lorenz curve. This was demonstrated by Kakwani (1993). I will briefly recap the necessary steps.

First, it is necessary to provide a formal definition of the Lorenz curve as can be found in Kakwani (1980), chapter 3:

$$L(p) = F_1(x). \tag{4}$$

The random variable x represents income, $L(p)$ is the Lorenz curve. The parameter p, the percentile of the distribution, is given by the income distribution function with

$$p = F(x) = \int_0^x f(t)dt.$$

The Lorenz curve $L(p)$ is $F_1(x)$, the first moment of the income distribution function $F(x)$, and it follows that

$$L(p) = F_1(F^{-1}(p)) = F_1(x) = \frac{1}{\mu}\int_0^x tf(t)dt,$$

with μ being the mean of x. The formula can be interpreted as follows: The parameter p gives the fraction of the population with incomes below x, while $F_1(x)$ gives the sum of all incomes smaller or equal than x as a share of total income. Graphically, the Lorenz curve is given by the set of all points $(F(x), F_1(x))$.

The slope of the Lorenz curve is given by $L'(p)$, with[1]

$$L'(p) = \frac{d(F_1 \circ F^{-1})}{dp} = \frac{dF_1}{dF} = \frac{x}{\mu}, \tag{5}$$

and the second derivative of $L(p)$ is given by

$$L''(p) = \frac{d^2F_1}{dF^2} = \frac{d}{dF}(\frac{dF_1}{dF}) = \frac{d}{dF}(\frac{x}{\mu}) = \frac{1}{\mu f(x)}. \tag{6}$$

[1]The derivation can be done step by step, first applying the chain rule and then the differentiation rule for inverse functions:

$$L'(p) = (F_1 \circ F^{-1}(p))' = F_1'(F^{-1}(p))(F^{-1})'(p) = F_1'(x)\frac{1}{F'(F^{-1}(p))} =$$

$$= F_1'(x)\frac{1}{F'(x)} = \frac{\frac{xf(x)}{\mu}}{f(x)} = \frac{x}{\mu}$$

The headcount index, the simple poverty measure we have used all along, depends on the poverty line z and the form of the income distribution $F(x)$, and it has a straightforward expression:

$$H \equiv F(z) = \int_0^z f(x)dx. \tag{7}$$

For deriving the GEP, we will need the slope of the Lorenz curve at the poverty line $L'(H)$:

$$L'(H) = L'(F(z)) = \frac{z}{\mu}. \tag{8}$$

Now, we take derivatives again with respect to μ and rearrange

$$\frac{\partial H}{\partial \mu} = \frac{-z}{\mu^2 L''(H)}. \tag{9}$$

This gives the derivative of the poverty headcount H with respect to μ, i.e. how poverty changes if mean income increases. By inserting equation (6) we get

$$\frac{\partial H}{\partial \mu} = \frac{-zf(z)}{\mu}. \tag{10}$$

Finally, multiplying with $\frac{\mu}{H}$ yields the formula of the the GEP:

$$\eta_H^\mu \equiv \frac{\partial H}{\partial \mu}\frac{\mu}{H} = \frac{-zf(z)}{H} = \frac{-zf(z)}{F(z)}. \tag{11}$$

This general formula can be used for calculations of poverty trends if we are willing to make assumptions on the exact form of the income distribution function $F(x)$.

Kakwani (1993) also showed how to derive an elasticity of poverty with respect to the Gini coefficient. What we are actually interested in is the effect of a change of inequality on poverty. But what exactly do we mean by a change of inequality? There are countless ways of how inequality, or in general, the form of the income distribution, can change. Kakwani proposed, as a simplifying measure, to just consider changes that correspond to a shift of the entire Lorenz curve. Such a shift directly corresponds to a change of the Gini coefficient and takes on the following form:

$$L^*(p) = L(p) - \lambda(p - L(p)). \tag{12}$$

where λ measures the percentage change of the Gini coefficient. If λ is positive, i.e. if inequality has increased, $L^*(p)$ is Lorenz dominated by $L(p)$. If we aggregate over the incomes in a society from the lowest income to any specified level x, this sum will always be lower in the situation characterised by the curve $L^*(p)$ than if the distribution is characterised by $L(p)$.

In addition, assuming that a change in inequality takes on this particular form also ensures that the income distribution before and after the change of inequality can be described by the same functional form. Of course, in practice this needs not be the case. So, all results based on this assumption need to be treated with care. If one were really interested in neatly decomposing the effect of a change in inequality and a change in mean incomes on the number of poor in a society, it is inevitable to use much more detailed information on the income distribution and its changes.

Kakwani (1993) derived the following formula for the Gini elasticity of poverty:

$$\eta_H^{Gini} = (\mu - z)\frac{f(z)}{F(z)} = \frac{z-\mu}{z}\eta_\mu^H. \tag{13}$$

Note that whenever $z < \mu$ i.e. the poverty line is below mean income, the Gini elasticity η_H^{Gini} is positive. Then an increase of the Gini coefficient will lead to an increase of the number of poor.

4.2 Fisk income distribution

The Fisk distribution, also known as log-logistic distribution is a two-parameter distribution. Compared to the Lognormal, it has heavier tails. It is fully described by the mean μ and the inequality parameter β. The latter is just the inverse of the Gini coefficient, thus $\beta = \frac{1}{G}$ with G standing for the Gini coefficient.

The cumulative distribution function of the Fisk distribution is given by:

$$F_{Fisk}(x) = \frac{x^\beta}{x^\beta + \alpha^\beta}, \tag{14}$$

where α is a scale parameter and given by:

$$\alpha = \frac{\mu}{\Gamma(1+G)\Gamma(1-G)}. \tag{15}$$

The density function of the Fisk distribution takes on the following form:

$$f_{Fisk}(x) = \frac{(\frac{\beta}{\alpha})(\frac{x}{\alpha})^{\beta-1}}{(1+(\frac{x}{\alpha})^\beta)^2}. \tag{16}$$

This information is all we need to calculate the GEP based on the Fisk income distribution:

$$\eta_{H,Fisk}^{\mu} \equiv \frac{-zf_{Fisk}(z)}{F_{Fisk}(z)} = \frac{-z\left(\frac{(\frac{\beta}{\alpha})(\frac{z}{\alpha})^{\beta-1}}{(1+(\frac{z}{\alpha})^\beta)^2}\right)}{\frac{z^\beta}{z^\beta+\alpha^\beta}}. \tag{17}$$

4.2.1 The Fisk vs. the Lognormal elasticity

Using GDP and Gini coefficients from the World Bank for the time span 1990-2012 and all countries included in the MDG database, I calculated both the Fisk elasticity and the Lognormal elasticity for the 1.9 US$ extreme poverty line. The data on income is GDP in 2011 PPP adjusted US$. The Gini coefficient is not readily available for all years. In between two observations of the Gini coefficient, I linearly interpolated any missing values. For the first and last years in the sample, I simply assumed that there had been no trend of inequality and replaced the missing values with the closest Gini observation. This approach is far from perfect and the dataset thus assembled should not be used for analyzing trends in inequality. Also it is not ideal for doing more in depth poverty analysis because, as argued before, GDP is the wrong income measure for doing so. But for the purpose of comparing the Fisk elasticity to the Lognormal elasticity, this dataset should be sufficient.

The absolute values of the Fisk elasticity are much lower than the Lognormal elasticity, and it displays a far lower variance. The mean unweighted value of the Fisk elasticity is -2.2 while the mean of the Lognormal elasticity is -3.8. When comparing the population weighted figures, both means increase slightly in absolute terms, to -2.5 and -3.9 respectively. The lowest value in absolute terms of the Fisk elasticity which I find is -0.09 and the highest value is -6.16, but the Lognormal elasticity has a far larger range, with a lowest value of -0.11 to a highest value of -27.68. Both elasticities increase (in absolute terms) with rising mean income and decrease with rising inequality. This is a result of their functional form, which could be derived analytically by taking derivatives. But as this is quite cumbersome, it might also be informative to perform a simple regression analysis to point out differences of the elasticities. A regression of the elasticity on the log of GDP gives a coefficient of -0.45 for the Fisk elasticity and of -2.01 for the Lognormal elasticity. A 10 percentage point increase of the Gini coefficient decreases the absolute value of the Fisk elasticity by 0.6 and the value of the Lognormal by 2.2.

In principle, while the Lognormal elasticity attains very high values especially when GDP is high, the Fisk elasticity is less responsive, and it is more in line with empirical values of the GEP. GDP per capita is far higher than household survey mean consumption but actually the latter is used for poverty anaylsis, and should also be used for calculating the elasticity. If the Lognormal elasticity is more responsive to mean income then the bias introduced by using the wrong means is much stronger for it than for the Fisk elasticity.

4.2.2 A short test of the formula

Bourguignon (2004) proposed a simple regression exercise to test the validity of the derived analytical expression for the elasticity. If the formula is indeed a good approximation of the "true" elasticity, in a regression of percentage changes of poverty on percentage changes of GDP times the analytical elasticity, the coefficient on the explanatory variable should be close to unity:

$$\frac{\Delta H}{H} = \eta_H^{\mu} \frac{\Delta GDP}{GDP} \text{ if } \eta_H^{\mu} = \frac{\Delta H}{H} \frac{GDP}{\Delta GDP}. \tag{18}$$

Bourguignon pointed out a problem with this simple approach. The measured percentage changes of poverty can become very large, especially when the poverty rate is very low. This simple test was repeated by Dellinger and Diprossimo (2015) with unbalanced panel data taken from the MDG data base in 2015. They first analyzed the annual percentage changes in the poverty headcount and found that their distribution is highly skewed with some huge increases of poverty, the maximum value being an increase of the order 4100%. Unfortunately, transforming all changes of the poverty rate into percentage changes entails the danger of magnifying measurement errors. The 4100% increase could for instance be explained if the poverty rate increases from 0.1% to 4.2%. Given the difficulties attached to measuring poverty, such fluctuations are frequent in the data. In order to prevent measurement errors from clouding the actual dynamics, Bourguignon (2004) proposes to restrict the sample to countries with a poverty rate above a certain threshold value at the beginning of the growth spell. Unfortunately, he does not specify further which threshold value should be used. Dellinger and Diprossimo (2015) thus chose a threshold of 5%.

Using the same data as Dellinger and Diprossimo (2015), I perform the same test for the Fisk elasticity and compare it to the result for the Lognormal elasticity. Both yield coefficients which are close to 1, i.e. 1 is included in the 95% confidence interval surrounding them. But the coefficient on the Lognormal elasticity (times GDP change) is 0.85 while the respective coefficient for the Fisk elasticity is 1.008. Thus, I find that the Fisk-elasticity is surprisingly close to a hypothesized "true" elasticity. Also, the coefficient of 0.85 on the Lognormal elasticity times the GDP change seems to indicate that the elasticity is too high.

4.2.3 Why the Fisk distribution?

Why should we choose the Fisk distribution to approximate the income distribution? Bresson (2009) concluded that, in general, 3-parameter distributions should be preferred. Yet, these distribution functions are far more complicated to work with, and interestingly, the Fisk distribution yielded elasticities which were in line with those derived from more complicated distribution functions (Bresson, 2009). To be more specific, Bresson (2009) tested a set of different distribution functions applied to data of the World Income Inequality Database (WIID) and, if feasible, calculated the subsequent elasticities. From initially 15 distribution functions proposed in the literature, 9 were chosen depending on whether they yielded Lorenz curves, which fulfilled some validity criteria. From these 9 distributions, GEPs for the 1 and 2 US$ poverty line were calculated. The mean and median of the GEPs for the 1 US$ poverty line are respectively -2.213 and -2.035. In absolute terms the Lognormal distribution yielded the highest elasticity with -3.52 and the Weibull elasticity was lowest with -1.18. The Fisk elasticity had a value of -2.11, thus close to both median and mean. A very similar pattern was observed for the 2 US$ poverty line. Again, in absolute terms the Weibull elasticity is lowest with -1.08, the Lognormal and the Beta 2 elasticities are highest with -2.71. With -1.88, the Fisk elasticity is close to the mean of all elasticities of -1.87.

Bresson went beyond such simple comparisons and tried to assess the different distribution functions based on their closeness to empirical Lorenz curves. He found that while the Fisk elasticity is clearly inferior to the 3 and 4 parameter distributions, of all the 2 parameter distributions it has the best goodness of fit. (This follows from the results presented in table 8.)

All of this makes me conclude that it is best to choose the Fisk distribution function for a cross country poverty analysis. First, the distribution function is easy to work with, as both cdf and pdf have closed forms and the only required data are mean income and the Gini coefficient, which are usually available. Second, and contrary to the Lognormal distribution, the derived elasticities are close to the mean of all elasticities. Third, the Fisk distribution function "predicts" a Lorenz curve which is on average closer to the empirical Lorenz curve than all other tested 2 parameter distributions.

5 Data: Survey means vs. NAS aggregates

For a sound GEP based poverty forecast, the growth rate of household consumption as measured by surveys is a most essential input. But there are no forecasts of household survey consumption, and publicly available data on past trends is scarce. The World Bank published internationally comparable household survey data on income and consumption, but only the time span from 2004 to 2012 is covered. For most countries in the dataset, two points in time are available, and so, for a very short period an annualized growth rate of survey consumption can be computed. This can be compared to the National accounts growth rate of household consumption.

The NAS consumption growth rate fluctuates a lot, probably due to some of the problems pointed out earlier. There are also some obvious outliers, e.g. in Georgia in 1996 consumption growth was 260%. In order to guarantee comparability, I took the NAS consumption figures for the years in which surveys were available and computed the annualized growth rate of NAS consumption analogously to the survey figures.

A first analysis of the data set on survey consumption provided by the World Bank showed that the average growth rate of survey consumption was 1.89%. For the same countries and the same time spans GDP growth equalled 2.69%, and NAS consumption growth equalled 2.55%. A regression of survey consumption growth on GDP growth (constant suppressed) yielded a coefficient of 0.69 for all countries. Although there are only 53 observations the coefficient is significant at the 1% level. When regressing survey consumption growth on NAS consumption growth suppressing the constant, the coefficient was 0.61.

Is GDP (or, for that matter, NAS consumption) growth an unbiased estimate of the growth of survey means? If yes, for a poverty forecast we could simply rely on the (available) forecasts of GDP growth. For this to hold true, the coefficient in the above regression should be statistically indifferent from 1.

We get to the following results: For the overall sample we can reject the null hypothesis that both growth rates are equal at the 1% significance level. But this seems to be mainly due to the surveys from Sub-Saharan Africa. They alone display a regression coefficient of 0.385 in the regression of survey consumption on GDP growth. If we exclude those, we cannot reject the null hypothesis any longer. The coefficient, without Sub-Saharan Africa, is 0.84. But then again, when excluding Africa, the results are driven by Eastern Europe and Central Asia, including many observations coming from rich countries, which are not interesting for poverty trend forecasts. Also, the sample is really small, and it is hard to draw convincing conclusions. Deaton (2005), discussing the results of Ravallion (2003b), argued that the essential fact of differential growth of survey consumption and NAS is lost if we do not take into account population size.

What would a population weighted model give?

A population weighted regression gives a much lower coefficient on GDP-growth, namely 0.58. It is statistically significantly different from zero and different from 1. The null hypothesis that the coefficient is equal to 1 can be rejected at the 0.1% significance level.

When using NAS consumption as explanatory variable, the coefficient β is a bit higher, equaling 0.64. It is statistically significant and it is possible to reject the null hypothesis that the coefficient is 1 at the 0.1% significance level.

Finally, for the population weighted regression, the results are not driven by the low ratio of Sub-Saharan Africa any more. Excluding Africa raises the coefficient (on GDP-growth) a bit to 0.66, but it continues to be significantly different from zero (at the 1% significance level).

Summary statistics on growth of GDP, growth of survey consumption and growth of NAS consumption are presented in table 2

Of course, my data only cover a very short time horizon and few countries. But they clearly indicate that NAS aggregates grow at higher rates than survey means. If this is indeed true, comparing my data with the results of past studies should show a growing discrepancy of the levels of NAS aggregates and survey means.

	$\frac{\Delta GDP}{GDP}$ all c. pop wei	$\frac{\Delta GDP}{GDP}$ all countr.	$\frac{\Delta survey}{survey}$	$\frac{\Delta GDP}{GDP}$ c. survey data	$\frac{\Delta NAScons}{NAScons}$ c. survey data
All	5.50	3.11	1.89	2.69	2.55
Sub-Saharan Africa	3.34	2.51	1.70	3.01	2.17
Europe and Cent. Asia	3.58	3.85	1.06	1.40	1.15
Latin America	2.59	2.60	1.82	2.68	3.79
N. Africa a. M.East	2.19	2.19	-	-	-
East Asia	8.29	4.27	3.85	4.47	3.15
South Asia	5.41	4.83	3.85	4.47	3.15

Table 2: A comparison of growth rates: GDP all countries population weighted, GDP all countries, survey consumption, GDP of only the countries with survey data and NAS consumption of countries with survey data

An interesting point of comparison is the work of Ravallion (2003b), who tried to answer a very similar question 15 years earlier, namely whether and to what extent NAS aggregates and household surveys agree.

Ravallion found discrepancies in levels, but when differentiating by whether a survey was an income or a consumption survey, some of the discrepancies vanished. For consumption surveys, the average ratio was 0.93 and the Null hypothesis of a ratio of 1 could not be rejected. But income surveys showed a ratio of 0.67, which was significantly different from 1. The data from the region Eastern Europe and Central Asia proved to be very problematic, showing only very weak correlation of NAS aggregates and survey means. One other region featured unusual results: India and Pakistan both had a ratio of survey means to NAS consumption of about 0.55, despite the surveys being consumption surveys.

The survey means published by the World Bank do not let me differentiate between income and consumption surveys. So, I have to build both the ratio of survey means to NAS consumption and to GDP. The income ratio will be downwards biased while the consumption ratio can be expected to be upward biased. In the overall sample, I find that the ratio of survey means to NAS consumption is 0.65. It is thus far lower than what was found by Ravallion (2003b) even though the estimate might be upward biased. This can be seen as an indication of different growth rates over the last 15 years. (Unfortunately, country coverage of both studies does not coincide and this is a reason for caution.) The ratio of surveys to GDP is 0.35, just half of the value found by Ravallion. This estimate is very low, but without knowing how many of the surveys actually are consumption surveys and how many are income surveys, we cannot make any strong statements. An analysis of the survey and NAS aggregates levels can be found in table 3. This comparison is a further indication that growth rates systematically differ.

6 Results

First, I will present a "reality check" of the poverty forecasting model. I analyzed the poverty trends in 6 countries. On average, the model performs quite well, but there are, of course, caveats. Most problems are related to the Gini elasticity of poverty in middle income countries. By construction the elasticity becomes very large and depending on whether inequality is falling or rising, poverty reduction is over- or underestimated.

Next, projections of poverty trends for the four regions with the highest poverty rates are presented.

6.1 Replicate past trends

Whenever I have data on household survey consumption, inequality and poverty for one country for at least two points in time, i.e. a growth spell, I can replicate the poverty trend in this growth spell by applying the GEP, $\eta^{\mu}_{H,t}$, and the IEP, $\eta^{Gini}_{H,t}$, according to equation (3). The mean of the distribution is given by the household survey mean.
Data on survey consumption means is available for a few countries and recent time periods on the World bank data base. For India and China, Deaton (2005) provides some data, which enables me to approximate the trend of survey consumption in the 1990s.

In order to test whether this simple model is capable of forecasting poverty trends if accurate data on household survey consumption is available, I replicate poverty trends in 6 countries.

6.1.1 China

China is extensively discussed by Deaton (2005). He concludes that NAS consumption growth and survey consumption growth diverge throughout the 1990s. While in 1990 the ratio of survey consumption to national accounts consumption was 0.95, it steadily declined only to reach 80% in 2000.

Private household consumption data from the national accounts in China is available on the world bank data base. In 1990, NAS consumption in China equalled 781.09 in 2011 purchasing power parity adjusted US$ per year. Applying the 0.95 ratio this gives daily survey consumption of 2.03297 US$, which is very close to the extreme poverty line of 1.9 2011 PPP US$. In fact, 60.7% of the Chinese population lived on an income below the extreme poverty line. But the next ten years were marked by an extraordinary growth performance of approximately 9.3% a year and a substantial reduction of poverty, which was nearly halved. In 1999, the extreme poverty rate was down to 36%.

The poverty count in China is based on consumption surveys, and thus for replicating the Chinese poverty trend, this is the necessary data. Deaton (2005) estimated that over the course of the 1990s, the growth of survey consumption was 1.7% lower than growth of NAS consumption. I calculated an approximate survey consumption series by taking the value of 2.03297 and letting it grow by the growth rate of NAS consumption minus 1.7 percentage points. In 2000, this approximate survey consumption equals 3.203 US$. Inequality increased in China during the 1990s with the Gini coefficient rising from 32.43 in 1990 to 39.23 in 1999. This is all the data necessary for calculating the analytical GEP and IEP, for both the Fisk distribution and the Lognormal distribution and to simulate the poverty trend in China in the 1990s.

The simulation for China is surprisingly accurate. Taking the Fisk elasticities, poverty is estimated to be at 35.6% in 1999, just 0.4% less than the actual value of 36%. This is a reassuring sign showing that the elasticities in combination with household survey consumption perform well in simulating poverty trends.

The same simulation was run with the GEP and IEP from the Lognormal distribution. Poverty was estimated to be 38.1%. Thus, I find that relying on the Lognormal distribution I derive a more cautious poverty trend forecast. Given the extraordinary performance of China, the Fisk elasticity actually performs better.

Of course, the development in China can hardly be considered as typical. Both the growth rate and the rate of poverty reduction were remarkable. If the poverty forecasting tool can only replicate extraordinary trends, it is of little use.

6.1.2 India

The poverty data for India reveals a rather unsteady pattern. In 1993 the official poverty count of the world bank is 49.4%. So, back in the early 1990s China was actually poorer than India. Then, for a long time, no data is available. Actually, a nation wide survey took place in 2000 and it showed a substantial decrease of poverty. But for reasons of survey design changes this survey was largely criticised and did not make it into the world bank poverty count (Deaton and Kozel, 2005). The next poverty data available comes from 2004, with poverty being 41.6%. This is actually a low decrease of poverty given the good growth performance of India in the 1990s and early 2000s with an average growth rate of 4.4% a year between 1994 and

2004. But from then on poverty reduction accelerates. In 2009, poverty is down to 32.7% and in 2011 it had reached 23.6%. GDP growth also accelerates to 6.3% a year, but not to the same degree as the acceleration of poverty reduction. A possible explanation for this trend could be that poverty reduction is measured in a very primitive way, and it only counts whether a person is below or above the poverty line. If people are far away from the poverty line, their incomes might grow without them being lifted above the poverty line. Also, it is possible that most growth takes place in the upper segments of the income distribution. If this were the case, inequality increases and this should be captured by an increase of the Gini coefficient. Unfortunately, measuring inequality is very difficult as was already discussed. Researchers have to rely on household surveys, which work well for the poor but badly for the rich (Székely and Hilgert, 1999). Indeed, for India there is some convincing evidence assembled by Banerjee and Piketty (2005) that income inequality is rising by more than shown in the official statistics. The rise of top incomes in India contributes to the widening gap between survey estimates of household consumption and NAS consumption.

For India two data sources were merged, namely data on the ratios of survey means to NAS consumption provided by Deaton (2005) for the 1990s and World Bank data on household survey consumption means from 2004 to 2011. India is one of the countries where household survey means and NAS estimates of consumption diverge most. For the household survey conducted in 1999/2000 the ratio had reached 0.56, an alarming low.[2] Indeed there is an inner Indian debate going on over how to reconcile surveys and national accounts (Deaton and Kozel, 2005).

The survey means of 2004 and 2011 provided by the World Bank are also far lower than NAS consumption, and the ratio of both has declined further. In 2004 the ratio was 0.59 and by 2011 it had decreased to 0.476. (From 2000 to 2004 an increase seems to have taken place, but it is not clear, whether both data series are indeed fully comparable, especially taking into account the problems with the 2000 survey.)

I constructed a series of survey household consumption based on the data provided by Deaton. I took 1999 as a starting value. NAS consumption per capita in 1999 was 1498.691, dividing by 365 and multiplying by 0.56 I get to my approximate value for daily survey consumption, 2.299 US$. According to Deaton (2005) NAS consumption grew on average by 1.1 percentage points more than survey consumption between 1983 and 1999/2000. I constructed survey consumption by shrinking the 1999 value by NAS consumption minus 1.1 percentage points. Unfortunately, this is a very rough

[2] Deaton (2003) analyzed the survey design changes of the 1999/2000 survey and proposed a method to adjust the results in order for them to be comparable to earlier surveys. The 0.56 ratio is calculated based on the thus adjusted data.

approximation, as I do not know any more details about surveys between 1983 and 1999. In 2004, survey consumption is available on the World Bank website. It is 2.814 US$. Between 1999 and 2004, I impute the yearly growth rate of survey consumption. Between 2004 and 2011, survey consumption growth was 3.696%.

Taking this series of survey consumption, I try to replicate the Indian poverty trend. It turns out to be far too optimistic. When I try to replicate the full trend from 1993 to 2011, my poverty estimate is 15.798% with the Fisk elasticities and 18.16% with the Lognormal elasticity. This is far from the actual value of 23.6%. I also tried to just forecast the development from 2004 to 2011. (This has the advantage that I do not have to rely on ratios and imputations, but that I can simply use the official data.) With the Fisk elasticities, I can replicate the trend from 2004 to 2011 quite well. My poverty forecast for 2011 is 23.94%. With the Lognormal elasticity the forecast is 25.62%.

A note of caution is necessary: The model turned low and stable inequality and strong growth in a poor country into a remarkable poverty reduction. But this might very well be the only thing the model can do; any more complicated developments cannot be captured.

6.1.3 Brazil and Ecuador - The middle income countries

So far, the model proved to be successful at replicating the trend in China, and at least parts of the trend in India could be replicated. Also, we have seen that so far, the Lognormal elasticities yielded the more cautious forecasts. Next, I will turn to two middle income countries, Brazil and Ecuador. The only source of the data is the World Bank; I do not have to do any imputations, but at the same time, only very short time spans are covered. In Brazil, data on survey consumption and poverty from 2007 and 2012 is available. Poverty in 2007 was just 5.8%, the Gini coefficient of 55.23 showed very high inequality, and the mean of household survey consumption equalled 13.986 US$. Until 2012 several favorable changes set in. Household consumption grew by a yearly rate of 4.54%, inequality fell with the Gini coefficient being reduced to 52.67, and poverty fell to 3.8%.

When trying to replicate the Brazilian trends, I found that I strongly overestimate the poverty reduction. My poverty forecast for Brazil is 2.07%, instead of 3.8%. A very important component contributing to the simulated trend is made up by the rate of change of the Gini coefficient multiplied by the Gini elasticity of poverty, $\frac{\Delta Gini}{Gini} * \eta_{H,t}^{Gini}$. This is due to the fact that by construction, the Gini elasticity of poverty becomes very large, if mean income μ and the poverty line z are far apart. Remember that $\eta_{H,t}^{Gini} = \eta_{H,t}^{\mu} * (\frac{z-\mu}{z})$. In Brazil, where μ was 13.986 in 2007 and growing, the difference to the poverty line z of 1.9 is large. While the Fisk growth elasticity of poverty in Brazil lies within the range of -1.68 and -1.83, the Gini elasticity of poverty lies in the range of 10.72 to 14.97.[3] Accordingly, the reduction of inequality gets a very strong weight within the trend forecast. Without it, poverty reduction would be 1.6 percentage points less.

In Ecuador the following data is available: Poverty in 2007 was at 6.8% and until 2012 it declined to 4%. Mean household survey consumption in 2007 amounted to 10.74, but its growth in the subsequent years was meagre, just amounting to 0.97% a year, but a substantial decrease of inequality took place. In 2007, the Gini coefficient was 54.33, and it decreased to 46.57 in just 5 years. This is equivalent to a $\frac{\Delta Gini}{Gini}$ of -3%. But measurement errors might also play a role.

Replicating the Ecuadorian poverty trend was impossible. The poverty reduction forecast is much too optimistic. Poverty in 2012 is estimated to be 1.38, instead of 4. The reduction is nearly entirely due to the trend component $\frac{\Delta Gini}{Gini} * \eta_{H,t}^{Gini}$. The Gini elasticity of poverty in Ecuador ranges from 7.7 to 10. Without this component, poverty would stall and only be reduced by 0.6%.

The examples of both Brazil and Ecuador clearly show the limits of forecasting poverty based on simple elasticities. Especially the Gini elasticity is not convincing, when the difference between the poverty line and mean of the distribution becomes large. This is not surprising as strong assumptions are needed for deriving $\eta_{H,t}^{Gini}$. Distributional changes can take on infinitely many forms, but by deriving the formula we restricted them to a single specific one. At the same time, many developments such as the poverty trend in Ecuador can only be explained by the changes of the distribution. Just taking out the term $\frac{\Delta Gini}{Gini} * \eta_{H,t}^{Gini}$ does not solve the problem.

What the forecasts also show is that redistribution in middle income countries, especially if it came in the form assumed by the model, could be very efficient in reducing poverty.

Finally, a note regarding the Lognormal elasticities. In line with the results of (Bresson, 2009) I find that the Lognormal elasticities are slightly larger in absolute terms than the Fisk elasticities in Brazil. In Ecuador, initially the Fisk elasticity is larger, but in the third year it is overtaken by the Lognormal elasticity. The poverty estimates hardly change when instead of the Fisk elasticity the Lognormal is used. They only become marginally more optimistic.

6.1.4 The African poverty trends - Congo, Rep. and Madagascar

The two African countries for which I chose to replicate trends are very different. While both are very poor by international standards, the Republic of Congo has a growing economy and it has witnessed strong poverty reduction. Madagascar is much poorer, and it has seen the economy falling, with household consumption falling and inequality rising.

[3] The elasticities grow over time (in absolute terms) as μ_t grows and inequality shrinks.

The key data for the Republic of Congo come from 2005 and 2011. Poverty in 2005 was 54.1%, with household consumption being 2.96 US$. Inequality was high with a Gini coefficient of 47.3. Household consumption in the period under consideration grew by 4.29% a year and the Gini coefficient fell to 40.2.

My simulation based on the Fisk elasticities shows poverty falling to 31.44%. The actual poverty count in 2011 was 32.8%. The simulation based on the Lognormal elasticities results in a poverty level of 34.3%. Given the huge decrease of poverty, both estimates are actually quite good, with the Fisk elasticities overstating poverty reduction and the Lognormal elasticities underestimating it. It also shows that for poor countries the forecasting method performs quite well and the approximations of the income distribution by known statistical functions are not too bad. The Lognormal elasticities actually give the more cautious approximation, and maybe, for forecasting purposes in poor countries, it should be preferred.

The data for Madagascar comes from the years 2005 and 2010. In 2005 82.4% of the population lived on less than the extreme poverty line of 1.9 US$ a day. Mean consumption was actually below the poverty line at 1.74 US$ a day. Not only was the country extremely poor, the situation even deteriorated. By 2010, household consumption had fallen to 1.45 US$ a day. Consumption growth was negative with a rate of -3.52%.

The poverty simulations show a terrible result with poverty rising to 91.7% (for both the Fisk and the Lognormal elasticities). The actual data show that poverty did increase, but by less than what the model would make us think: The poverty level in 2010 in Madagascar was 87.7%. The most recent history of Madagascar is marked by political instability. In 2009 protests against the president led to a military coup followed by international sanctions and a political crisis, which lasted until 2011. Thus the household survey from 2010 took place during a period of political turmoil, which might cast doubt on their quality.

6.1.5 Lessons to be learnt

For forecasting poverty trends, several conclusions can be drawn.

1) The model performs quite well for good growth spells in poor countries, when there is not a lot of change in inequality. This is what we need the most. It bears the risk of overestimating poverty reduction and producing overly optimistic forecasts. A possible safeguard could be to try out different elasticities (Fisk, Lognormal, or others) and stick to the more conservative forecasts.

2) The model does not do a good job of forecasting poverty when there are large changes of inequality in middle income countries. For forecasting, this is not so problematic as, so far, no consensus regarding future trends of inequality has been reached. Usually, the baseline assumption is no change

of inequality.[4]

3) For understanding poverty we need good data on household consumption of the poor. For forecasting poverty trends, we need to know about survey consumption levels and growth, and, so far, the publicly available data is clearly insufficient.

4) Even though the model is simple, it seems to be a handy tool for quick forecasts. Also, it might help to get a feeling for the plausibility of episodes of poverty reduction, and when to expect severe data measurement errors either in terms of the growth rates or in terms of the inequality measurement.

6.2 Forecasts

With the Sustainable Development Goals (SDGs), the international community has set out yet another ambitious goal - namely that poverty be eradicated. In 2030, nobody should be left with an income below the extreme poverty line - except for exceptional circumstances like in the aftermath of a natural disaster.

What can the model described by equation (3) tell us with regards to the SDGs? The UN presented figures for 2011 on extreme poverty in the different world regions, which document the early achievement of the first target of the Millennium Development Goals (MDGs). That is, on a global scale the percentage of people living in extreme poverty was cut by more than half four years ahead of time (UN-MDG Report, 2015). I took the UN data on poverty counts as a basis for regional poverty estimates.
Unfortunately, all other input data required are not readily available. I took inequality data from the data set used for comparing the Fisk and Lognormal elasticity in section 4.2 and formed the population weighted average of the Gini coefficient. For μ, I had to resort to an analysis of the insufficient data on survey consumption. I have data on daily survey consumption for 64 countries and among those, 53 provide data from two different years, allowing the calculation of a growth rate. For those 53 countries, I linearly interpolated survey consumption for the years within the growth spell. If the last observation on survey consumption comes from the years 2009 or 2010, I updated survey consumption using the GDP growth rate as a basis. (Of course, this is not perfect, but it still seemed to be the most straightforward approach.) For all other countries, I proceeded as follows: I calculated the regression ratio of survey consumption to national accounts surveys, (suppressing the constant). This gives a coefficient of 0.65, (and a reassuring R^2 of over 0.94). For all countries which did not have data on survey consumption, I took survey consumption to be best approximated by NAS consumption * 0.65. Thus, I assume that the ratio of survey consumption

[4] Deininger and Squire (1998) have shown that the Gini coefficient does not systematically vary with growth and that income distributions, (at least to the extent that we are able to measure them), are actually quite stable.

Country	Year	Poverty	S.Cons	Gini	Pov_{Fisk}	η^{μ}_{Fisk}	η^{Gini}_{Fisk}	Pov_{Lnorm}	η^{μ}_{Lnorm}	η^{Gini}_{Lnorm}
Brazil	2007	5.8 %	13.99 $	55.23		-1.68	10.72		-1.67	10.60
	2012	3.8 %	17.46 $	52.67	2.07 %	-1.83	14.97	1.97 %	-2.06	16.91
China	1990	60.7%	2.03$	32.43		-1.28	0.09		-1.16	0.08
	1999	36.0%	3.03$	39.23	35.62%	-1.60	0.94	38.15%	-1.36	0.81
Congo, Rep.	2005	54.1 %	2.96 $	47.32		-1.11	0.62		-0.93	0.52
	2011	32.8%	3.81 $	40.17	31.44 %	-1.84	1.85	34.32 %	-1.59	1.60
Ecuador	2007	6.9 %	10.74 $	54.33		-1.65	7.7		-1.54	7.17
	2012	4 %	11.27 $	46.57	1.38 %	-2.04	10.09	1.36 %	-2.12	10.34
India 93	1993	49.4 %	1.96 $	31.00		-1.27	0.04		-1.17	0.03
	2004	41.6 %	2.81 $	33.15	27.45 %	-1.96	0.94	29.49 %	-1.70	0.82
	2011	23.6 %	3.63 $	33.75	15.80 %	-2.35	2.13	18.16 %	-2.14	1.95
India 04	2004	41.6 %	2.81 $	33.15		-1.96	0.94		-1.70	0.82
	2011	23.6 %	3.63 $	33.75	23.94 %	-2.35	2.13	25.62 %	-2.14	1.95
Madagas-car	2005	82.4 %	1.74 $	38.88		-0.74	-0.06		-0.72	-0.06
	2010	87.7 %	1.45 $	40.63	91.67 %	-0.50	-0.12	91.69 %	-0.52	-0.12

Table 4: Poverty trends replicated

to NAS consumption for countries, for which I do not have data, is the same as for those for which I have some data. In 2011, survey consumption in the countries with data equalled 12.8. For the same countries, in 2011 the average of NAS consumption * 0.65 equals 12.3. Thus, maybe I risk underestimating the level of household consumption in 2011. Summary statistics on survey consumption, GDP, NAS consumption and the Gini coefficient are presented in table 3.

Region	μ	GDP	NAS cons	Gini
Europe and Cent. Asia	17.35	49.96	26.75	34.67
N. Africa a. M. East	12.70	32.41	16.91	34.92
South Asia	3.68	12.21	7.46	33.16
East Asia	6.78	26.82	10.21	37.71
Latin America	14.61	37.03	21.44	49.91
Sub-Saharan Africa	3.27	9.10	5.51	42.87

Table 3: A comparison of household survey consumption, μ, GDP and NAS consumption per day and the Gini coefficient in 2011

More difficult than finding a starting value for μ is to settle for convincing growth rates of survey consumption. As was shown before and in line with the results of Deaton (2005), it is not correct to simply assume that survey consumption grows at the same rate as GDP.

For the following projections I will mostly analyse what would happen if the rate of survey consumption growth recorded in the late 2000s remained the same. I resort to the population weighted averages of the growth rate of survey consumption for four important regions, East Asia, South Asia, Sub-Saharan Africa and Latin America. For each country, I only have 1 observation on annualized survey consumption growth stemming from one growth spell of on average 5 years.

I decided not to do any projection for the region North Africa and the Middle East. First, there are just 3 observations on survey consumption. Secondly and more importantly, out of the 11 countries with poverty data in the region, 3 have a civil war going on (Syria, Yemen and Iraq), one is occupied (West Bank and Gaza) and one experiences mass immigration by refugees fleeing the before mentioned civil conflicts (Jordan), not to mention the political instability in Egypt and other countries in the aftermath of the Arab Spring. A proper analysis of future poverty trends in the region requires much more than simply looking at growth rates and inequality trends.

6.2.1 East Asia

Poverty in East Asia has declined dramatically in the last two decades. This is mainly due to the tremendous poverty reduction in China, but also the other countries in the region were very successful. In 2011, the extreme poverty rate in the region was 12%, according to the UN-MDG Report (2015). Inequality in the region has increased, but to what extent exactly is disputed. I took the Gini coefficient to be 37.71 in 2011 which is the population weighted mean of the Gini coefficients of East Asia reported in the World Bank database.

On survey consumption growth, I have 6 observations for the 11 countries in the region (with some sort of poverty data). For 3 out of these 6, no corresponding survey consumption levels are available, and so it is not clear what is the reference time period, (i.e. the growth spell), and they were not part of the previous analysis. But fortunately, the 3 "new" observations now include China, (with a survey consumption growth rate of 7.86). The population weighted average of survey consumption growth for these 6 observations is 7.19%.

But the second largest country in the region and fourth largest in the world, Indonesia, does not have any consumption survey data. The poverty rate of Indonesia in 2011 was 16.2%, and with a population of 244 million (in 2011), this gives about 39.5 million people below the extreme poverty line. Annualized growth in Indonesia between 2006 and 2012 was 4.5%, far below the population weighted average of survey consumption growth in the region. An analysis of survey consumption growth and GDP growth as in section 5 shows that for Asia, survey consumption is best predicted by GDP growth times a coefficient of 0.8. (As this is in line with the results of other authors dealing with the same problem, e.g. Chandy et al. (2013), I will take the result at face value, although it is based on very few observations.) Thus, my best guess for the growth rate of survey consumption in Indonesia is 3.6%.

Including the guess for Indonesia, the population weighted average of survey consumption growth becomes 6.71%. I take this value for my baseline growth rate for the projection of poverty trends in East Asia until 2030. But I also consider a pessimistic trend with growth of only 4%. Long-term growth forecasts for China of the Economist Intelligence Unit (EIU) show a decline of GDP growth from 9.5% in 2011 to 4.7% in 2020. Thus I deemed it justified to consider a "low" growth scenario as well. Also, in terms of sensitivity analysis, I considered inequality trends of +/-0.5% growth of the Gini coefficient per year for both growth rates. A decrease of the Gini coefficient by 0.5% a year results in a Gini coefficient of 34 in 2030, while an increase results in a Gini coefficient of 41 in 2030. Thus, I do not assume any strong trends of inequality, but just minor variations.

Figure 1 shows the results of the projections. Until 2020 we can expect

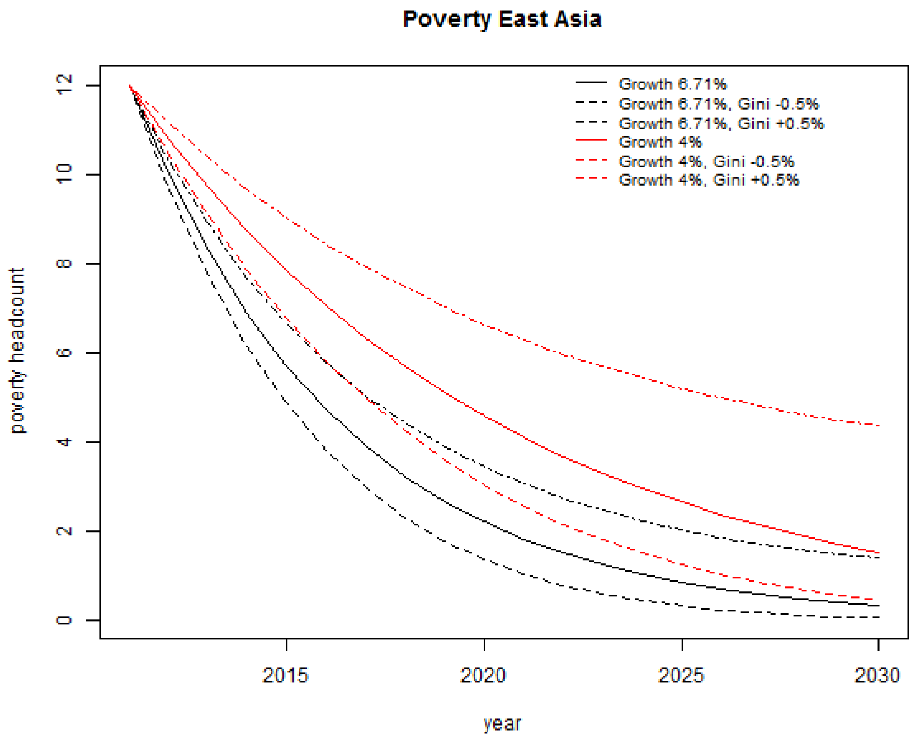

Figure 1: Poverty projections for East Asia

a rapid decline of poverty to 2.2% in the baseline scenario of strong consumption growth of 6.71% a year and no changes of inequality (the solid black line in figure 1). Until 2030, poverty will be far below 1%, (the exact forecast being 0.3% for 2030). Yet, we see that these results are sensitive with respect to future inequality. Unfortunately, the development of China has been characterized by an increase of inequality especially in the 1990s. Official numbers for the 2000s show a levelling off of inequality, and even a stark decline from 2010 to 2011. These numbers have been disputed by many independent researchers, who claim that inequality is much higher, with the Gini coefficient above 50 (Xie and Zhou, 2014). If we assume a slow rise in inequality, poverty only declines to 1.4% (the upper black dashed and dotted line).

The red line represents the pessimistic growth scenario with consumption growth of 4%. Without any changes of inequality poverty would decline to 1.5%. Declining inequality and low growth, the lower red dashed line, would be as effective as stable inequality and high growth in achieving poverty reduction. The only truly pessimistic outcome, low growth and increasing inequality represented by the upper red dashed and dotted line shows poverty at 4.36% in 2030.

A note of caution regarding the results is necessary: The concept of elasticity makes it close to impossible to actually reduce poverty to zero. When poverty is at 1%, the elasticity is -2 and income grows by 10%, poverty will decline by just 0.2 percentage points to 0.8%. On the one hand, this seems overly pessimistic. On the other hand, it is not so unlikely that there is a small fragment of the population who simply does not profit from economic growth. (Imagine old, disabled or sick people. Their consumption does not depend on the situation on the labour market, but rather on the social assistance/protection system or the willingness of their families to support them. Also indigenous groups who preserve their original lifestyle are usually cut off from economic development.) In many Latin American countries incomes have been relatively high for some time now, but the low poverty rates have not shown much change.

We can conclude that unless inequality severely worsens, the sustainable development goals are well within reach for East Asia.

6.2.2 South Asia

In 2011, the UN reports poverty in South Asia at 23%. Will it be possible to reduce poverty to close to zero until 2030? Inequality is quite low, with a Gini coefficient of 33.16, but so are incomes. μ in 2011 equals 3.68, less than twice the extreme poverty line. In fact, in terms of average incomes South Asia is not so different from Sub-Saharan Africa.

For South Asia, six observations on survey consumption growth were available including India. (No survey consumption levels are available for Bangladesh, but a growth rate is provided.) The coverage in terms of survey data for South Asia is really good. I have 6 observations for 7 countries with poverty data. Only the Maldives are missing (and of course countries like Afghanistan, who do not even provide poverty data.) The population weighted average growth rate was 3.34%, which I take to be the baseline growth rate. I also consider a more optimistic growth rate of 4.32. This is the result of applying the coefficient of survey consumption growth to GDP growth for Asia of 0.8 to the population weighted GDP growth of South Asia, which is 5.4.

Also trends in inequality of +/-0.5% are considered. They correspond to either a decline of the Gini coefficient to 30.15 or an increase to 36.46 until 2030.

Figure 2 displays the results for South Asia. The solid black line gives the baseline growth path with stable inequality. Under these conditions poverty will decline to 3.75% until 2030. The optimistic growth scenario results in a poverty level of 2.06%.

With a Gini coefficient of 33, inequality is rather low in South Asia. At the same time, household consumption (at least in 2011) is very low. Thus,

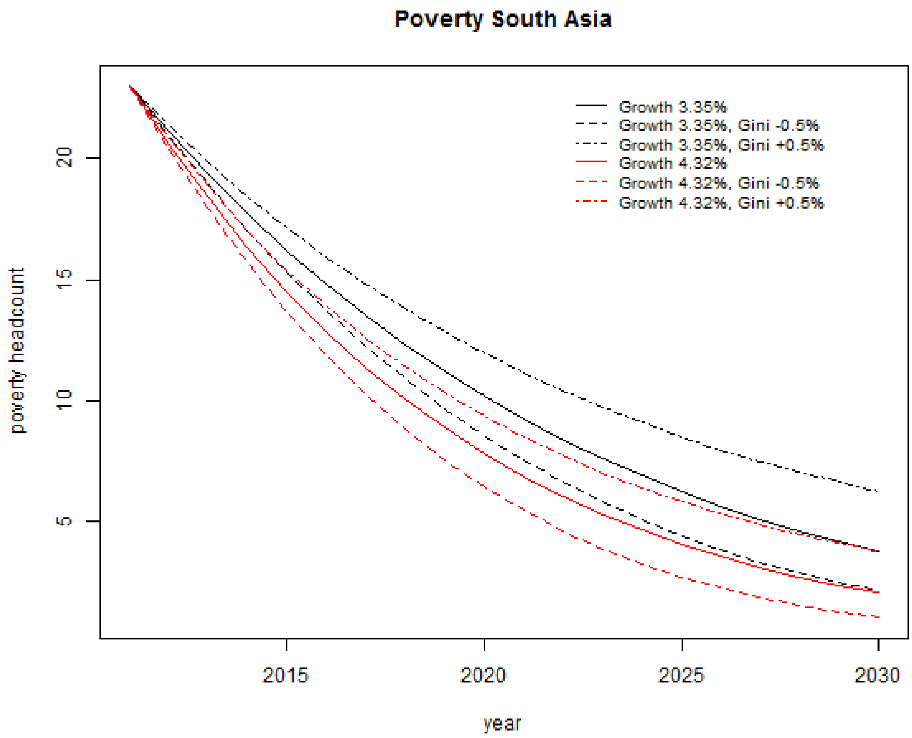

Figure 2: Poverty projections for South Asia

the inequality elasticity of poverty is also rather low and the poverty trends assuming inequality changes do not deviate much from the baseline trends. Until 2030 a one percentage point higher growth rate has about the same effect as a 0.5% per year decrease of inequality.

In the best case scenario of high growth and declining inequality, poverty is projected to be 1.02%, while in the worst case scenario of low growth and increasing inequality, it is 6.23%.

To sum up, it is perfectly within reach to reduce poverty in East Asia to well below 5%. But for poverty to come down below 1%, an increase of growth with respect to the status quo maybe in combination with a decline of inequality are necessary. Accompanying policy measures e.g. establishing a social safety net could alter the picture.

6.2.3 Sub-Saharan Africa

Sub-Saharan Africa is the only region where the first Millennium Development Goal was not met. Poverty in 2011 is still at 47% according to the UN-MDG Report (2015). Thus, it is clear that achieving the first SDG-target will also be hardest in Sub-Saharan Africa. Nevertheless there is reason for hope. In recent years many countries could greatly improve their growth performance, and also in the areas of health care and schooling important progress has been made (McKay, 2013).

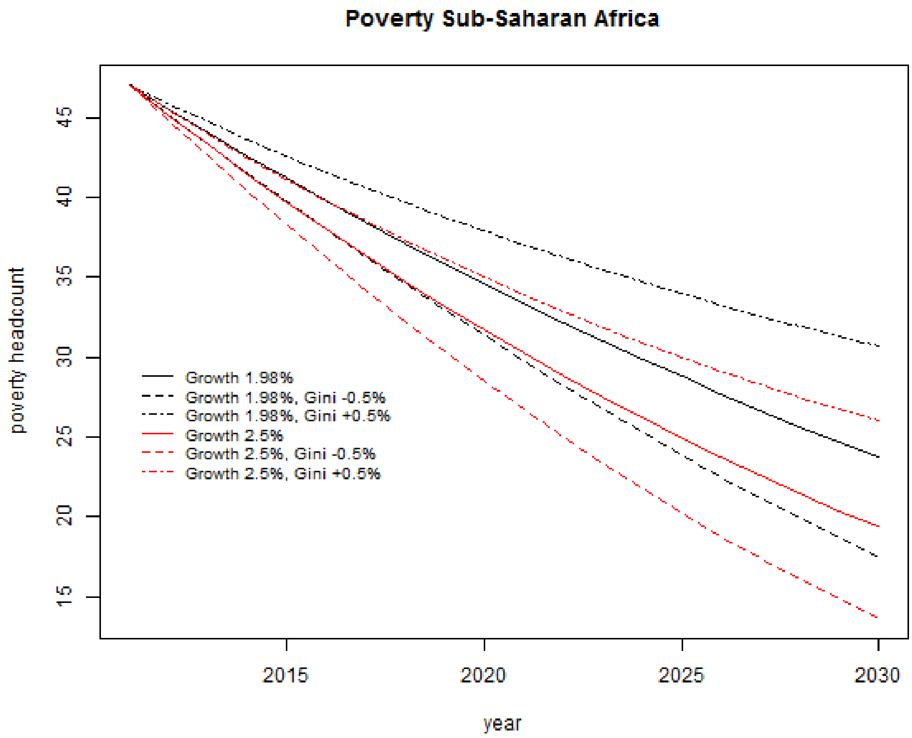

Figure 3: Poverty projections for Sub-Saharan Africa

Unfortunately the World Bank survey data shows that GDP growth and survey consumption growth in Sub Saharan Africa are quite detached. The OLS regression suppressing the constant of survey consumption growth on the corresponding GDP growth presented in section 5 shows an insignificant coefficient of 0.385. While we have only very few observations, 14 in total, it has to be pointed out that for the other regions, (Latin America, Europe and Central Asia and South and East Asia combined) the number of observations is similar but the results are significant and the coefficients are considerably higher. Maybe this is a result of the importance of subsistence farming in Sub-Saharan Africa. It can be expected that for people depending on subsistence farming, the effect of growth in the urban and formal economy will be negligible for their consumption opportunities. Also, the quality of national accounts in Sub-Saharan Africa has been challenged (Jerven, 2009).

Survey consumption growth in the 14 Sub-Saharan African countries for which data is available averaged 1.7%. Weighted by population the average becomes 1.98%. In the end, I chose to consider the following growth rates: 1.98% as representing a continuation of the current growth profile and 2.5% as an optimistic scenario. The average GDP growth rate was 2.51% and weighted by population size it became 3.34%. A survey growth rate of 2.5% is optimistic but not unrealistic. As always, Gini trends of +/-0.5% were considered.

Figure 3 shows the results. The continuation of the current trend is represented by the solid black line. It shows a poverty decline by a little less than half and a poverty level of 24.89% in 2030. This is not nearly enough to reach the first target of the SDGs. A decline of inequality and a more optimistic growth outcome, the red dashed line, shows poverty decreasing to a level of 15.5%.

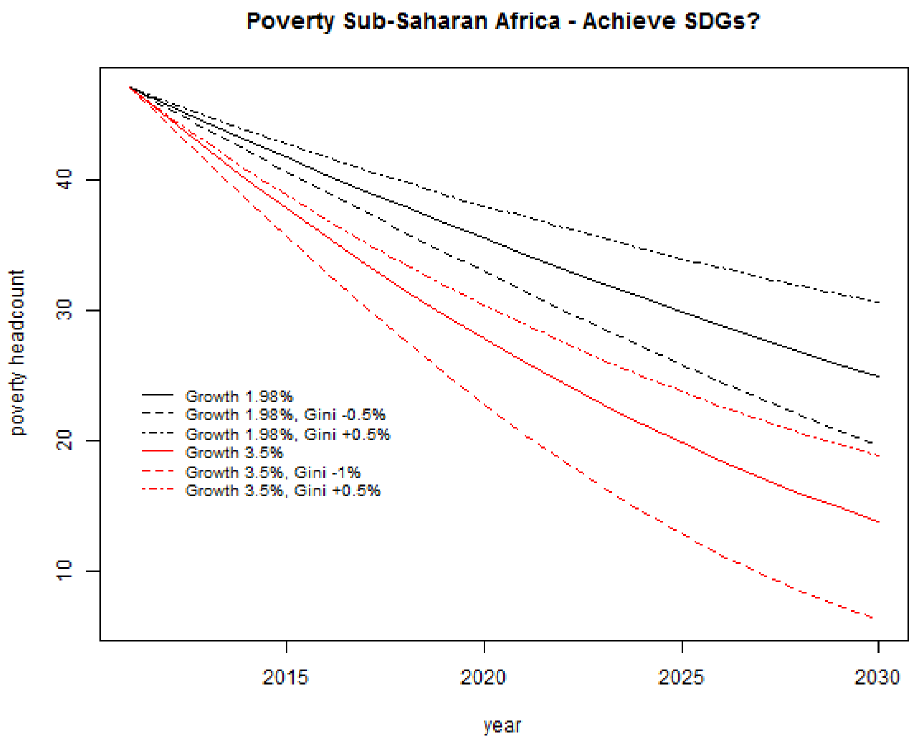

Figure 4: Sufficiently optimistic assumptions to lower the poverty projections for Sub-Saharan Africa to 5%

What would be necessary to actually achieve the SDGs? Obviously, growth has to speed up. But in Sub-Saharan Africa, there is also considerable scope for reducing inequality. In 2011, the Gini coefficient was 42.9. If it decreased by a rate of 1% per year from 2011 to 2030, the result would be a Gini coefficient of 35.44 at the end of the period. This does not imply an impossible inequality trend - the Gini coefficient of a country can fluctuate quite a lot - and it would still be higher than the Gini coefficient of South Asia today.

If survey consumption growth increased to 3.5% a year and inequality decreased by 1% a year, the poverty rate in Sub-Saharan Africa could actually be lowered to 6.3%. The results of the simulation are displayed in figure 4 with the red dashed line showing the potential path to substantially reducing poverty.

A poverty level of 6.3% is not equivalent to achieving the goals, but it is far better than what a simple continuation of the status quo would give. Also, the low inequality, high growth path only results in an increase of μ from 3.27 to 6.29 US$ a day (adjusted for 2011 purchasing power parity). This leaves a lot of room for income growth.

At the moment, it seems possible to substantially reduce poverty in Africa in the next 15 years or so, but if poverty should really be eliminated, additional and targeted policy measures, (maybe financed by the international community?) will be necessary.

6.2.4 Latin America

Poverty rates in Latin America have been low for a while now, but extreme poverty has not yet disappeared. According to the UN, in 2011, 5% of the population still lived in extreme poverty. Incomes are much higher than in Asia, with μ being 14.56 in 2011, but the region is characterized by the highest inequality of all. The Gini coefficient is 49.91.

Data on survey consumption growth is available for 15 countries, out of 25 with data on poverty. The population weighted average of survey consumption growth was 2.56%, nearly the same as the average GDP growth rate of 2.60%, or its population weighted version, which is 2.58%. I do not find the usual discrepancy between survey consumption growth and GDP growth. Only the simple, not population weighted average of survey consumption growth is substantially lower with 1.82%.

For the projection I took 2.56% to be the baseline growth rate, and I analysed the usual inequality trends of +/-0.5%. Figure 5 shows the results. Immediately, it becomes apparent that the real challenge is not growth but inequality reduction. At the same time, it helps to recall the problems with my forecasting model based on equation (3). When μ is far higher than the poverty line, the inequality elasticity becomes huge, and even minor trends in inequality become crucial for the poverty trend.

The baseline scenario with no changes of inequality sees poverty gradually decline to 1.89%. Thus growth alone at the same speed as today seems to be insufficient to reach SDG-target 1. If inequality were to decrease, the picture is changed dramatically. Poverty decreases to 0.30%.[5] But if

[5] Due to the multiplicative structure of the model, poverty cannot decrease to zero.

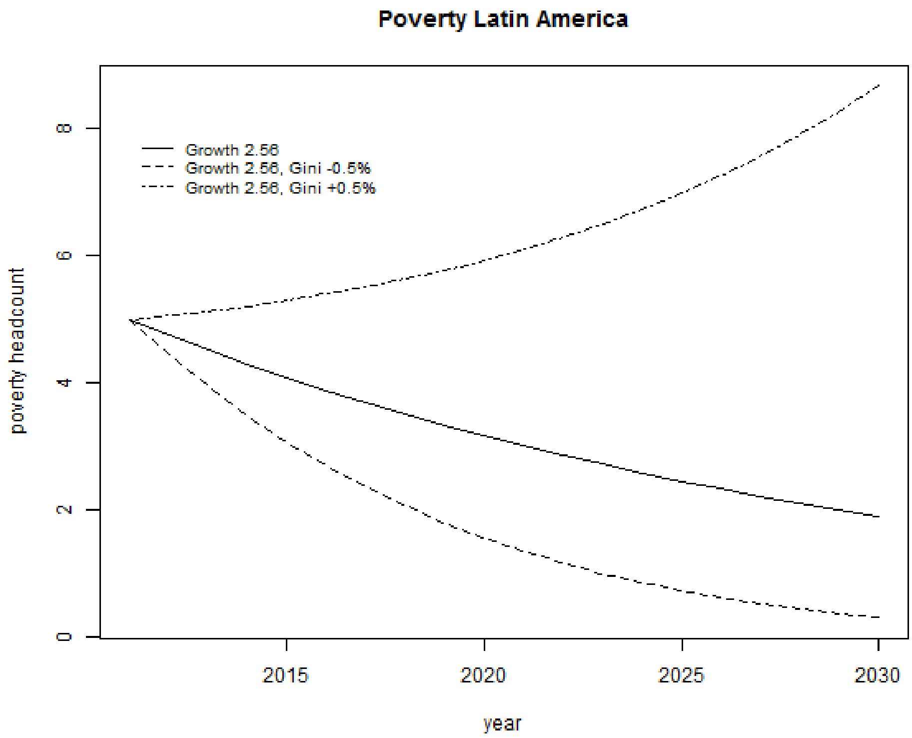

Figure 5: Poverty projections for Latin America

inequality were to increase, poverty would actually increase as well, even though there is positive income growth. The poverty rate in 2030 would be 8.68%.

Interestingly, in the last decade or so, many left wing governments came to power in Latin America. They tried to establish social assistance programs for the poor (for instance in Brazil, by the government of Lula da Silva) and to redistribute the gains from natural resources exploitation. Actually, between 2000 and 2011 the Gini coefficient in Latin America has declined by 5.2 percentage points from 55.1% to 49.9%. Right now, in 2016 we see a backlash against some of these governments, again in Brazil but also in Venezuela, (although the situations are not alike.) It will be interesting to see whether and how this will impact inequality in Latin America.

7 Discussion

What do these results imply for reaching target 1 of the Sustainable Development Goals? In short, a continuation of the 2000s growth trend with no changes of inequality will get us quite far in both East Asia and Latin

America. But adverse inequality trends pose a threat to these outcomes, particularly so in Latin America. A continuation of the current growth trends in Sub-Saharan Africa is not enough to achieve the target. Indeed, getting from 47% poverty to close to zero would require truly tremendous progress. I estimate that current trends would see Africa approximately halving its poverty rate - getting it down to 24.8%. In South Asia, a continuation of current trends is expected to lower poverty to below 4%. While this is clearly better than the situation in Sub-Saharan Africa, it is not exactly what people understand by ending poverty.

Poverty reductions are modeled as a non-linear process, especially when poverty rates are very low. And in middle income countries, poverty outcomes are increasingly sensitive to distributional changes.

Other researchers have come to the same overall conclusion, namely that good progress against poverty is to be expected but getting down to virtually zero is unlikely. I will shortly review the main results of two other studies who tried to forecast poverty trends until 2030. Ravallion (2013) analyzed growth and poverty reduction trends in order to analyze how long it will take to lift those remaining poor out of poverty. At the turn of the century there was a break in growth trends according to the household surveys data. While growth of household surveys had been 0.9% per annum in the 1990s, in the 2000s it increased to 4.3%. This is largely due to the catch up process most developing countries debarked on, and it cannot be attributed to China alone. The increase of growth without major setbacks in terms of inequality for the developing world as a whole resulted in a practically linear decline of poverty since 1990. On average, poverty declined by 1 percentage point per year since then. While in the 1990s most of the poverty reduction was due to the progress made by China, other regions are catching up now and seeing notable progress. Outside of China the poverty rate declined by 0.4 percentage points in the 1990s but by 1 percentage point since the turn of the century. Of course, a linear projection of this positive trend suggests that poverty can indeed be brought down to very low levels by 2030. In the year 2027, the poverty rate in the developing world would be at 3%. If instead growth and poverty revert back to their 1990s trends, it would require another 50 years to eliminate poverty (Ravallion, 2013).

Ravallion (2013) notes that the observable linear trend of poverty is indeed the product of nonlinear driving forces. The underlying GEP must have increased and growth must have accelerated. Nevertheless, it is most unlikely that the trend will indeed continue to be linear for long. Especially when the poverty rate drops below 10% the nonlinear character of poverty reduction processes is expected to become apparent. One way to deal with this, when forecasting poverty, is to assume stability of the income distribution and to project the distribution into the future. Based on this simulation approach, Ravallion (2013) calculates that an annual growth rate of 4.5% will be necessary if the poverty rate should drop to 3% by 2027. Hereby the author assumes that global inequality will remain at 2008 levels. If inequality were to drop to the 1999 level, only 3.4% growth would be necessary to achieve the same target. A growth rate of 4.5% of household consumption as measured by surveys is not easy to achieve - especially as the contribution of China is increasingly irrelevant, as extreme poverty in China is being eliminated. The most recent household surveys in Sub-Saharan Africa found consumption growth to be nearly 2% a year, less than half of what would be necessary.

Why should poverty trends eventually become nonlinear? A change of the poverty rate is the direct result of the interaction of the shape of the income distribution, its position with respect to the poverty line and income growth. Growth has the largest impact on the number of poor, if the mode of the income distribution is at the poverty line (Chandy et al., 2013).

Chandy et al. (2013) analyzed the past poverty reductions by focusing on exactly this interaction. They focus on the number of people with a daily income of 1.20 and 1.25 US$ a day, i.e., people who are likely to cross the poverty line within the next year if their incomes grow. The number of these "immediate" poor was a little over 100 million in the 1990s, which enabled high poverty reduction with relatively little growth. By 2010 this number had declined to 85 million people. But stronger growth enabled an equally rapid poverty reduction. In the future, this number will decline further, to 56 million in 2020 and 28 million in 2030 according to the baseline scenario of Chandy et al. (2013). The authors emphasize that poverty reduction can also be decomposed by regions. While in the 1990s it was mainly driven by China, now most of the poor situated directly below the extreme poverty line live in India. Currently, the mode of the Indian income distribution is situated at the poverty line and good progress can be expected. But in Sub-Saharan Africa the mode of the income distribution is still far below the poverty line. This slows down progress against poverty. It will need a tremendous effort to move out of this deep poverty.

Chandy et al. (2013) also offer some other reasons, why there are diminishing returns of equitable growth for poverty reduction. First, as countries become richer, poverty is increasingly concentrated in so-called "pockets of poverty". These might be geographically isolated populations, or minorities who face discrimination and are therefore unable to share in income growth. Also, poverty will increasingly become a problem of marginalized people, i.e. the sick, the disabled, the elderly and orphans. Instead of economic growth a social safety net, be it provided by the families or the state, is necessary for helping them.

At the same time, poverty will become concentrated in fragile and failed states. Currently one third of the poor in the world live in fragile states, but by 2030 this fraction is expected to increase to two-thirds.

The overall conclusion of Chandy et al. (2013) is that while the goal of ending extreme poverty by 2030 is a noble one, it will be extraordinarily hard to reach, and it will require many initiatives at the local level targeting the especially vulnerable. It is not impossible to end poverty, but only relying on a continuation of the status quo will not be enough.

As for my own results, I find it reassuring that they are broadly in line with these other studies, although the methods employed differ. For instance, Chandy et al. (2013) project poverty in Sub-Saharan Africa to be 23.6% in 2030 according to their baseline scenario. This is below but still close to my baseline result of 24.9%.[6]

8 Conclusion

In this paper the starting point was the debate on the growth elasticity of poverty, the GEP. Different estimation and derivation methods have been proposed in the literature, which yield markedly different results. But many of these differences become easily explicable, if we take into account the important insights that were gained concerning the measurement of poverty. Most importantly, it makes a huge difference whether income growth and mean income is measured within the National Accounting System (NAS) or via household surveys. Then I proceed by proposing an alternative analytical elasticity. Instead of the widely used Lognormal elasticity I propose a Fisk elasticity, and I explore some of its features. I propose a simple model for analysing poverty trends based on both the GEP and the IEP, the inequality elasticity of poverty. This model is first used to replicate past poverty trends, in order to check whether it is at all useful. Finally, poverty forecasts are made in order to assess whether the first target of the SDGs, eliminating poverty until 2030, can be reached. The forecasts are based on household survey data and are done separately for four world regions, East Asia, South Asia, Sub-Saharan Africa and Latin America. Finally these results are compared to other forecasts of poverty trends.

What lessons can be drawn from the study? The lessons are threefold. First, it has been shown that the GEP is actually a useful tool. A lot of the controversies in the literature questioning the GEP can be traced back to confusion about different data sources. Especially the analytically derived elasticities prove to be useful for performing a fast-track simulation of future trends, if the correct data is fed into the model. I proposed the use of the Fisk elasticity as an alternative to the Lognormal elasticity. The Fisk elasticity is more robust to misspecified means, which became apparent when calculating both elasticities for all country-years in the MDG database using GDP per capita as mean income. But when the correct household survey means are used, the difference between both elasticities actually becomes very small. The forecasts are also based on the IEP, the inequality elasticity of poverty. In order to analytically derive the IEP, very strong assumptions need to be

[6]Both studies do not offer any other projections of poverty by region.

made about the exact form of a distributional change. Of course, in practice these assumptions do not hold. Misspecification of the IEP are especially relevant, if the distance between mean income and the poverty line is large. Thus replicating and forecasting poverty trends in middle income countries is difficult.

Second, the work raises questions regarding our economic welfare measures. Replacing GDP by household survey consumption means has proven crucial for our results. We are used to concentrating on GDP growth and to, for instance, assess a government's performance against this benchmark. But if a lot of the growth that is measured does not actually reach households, or at least not those who are situated more towards the bottom of the income distribution, the focus on GDP growth might be ill placed. For researchers and the interested public it matters to know what GDP and other NAS aggregates actually measure and how reliable these figures are, in developing countries and elsewhere. Knowing about the caveats of conventional measures will help interpreting research results. In this respect the Gini coefficient is especially problematic. If country differences in the level of inequality are primarily reflecting differences in survey design, there is no point in even running cross country comparative studies of inequality. And, of course, neither GDP nor NAS consumption nor household survey consumption are actually measuring well-being.

Third, the results of the trend simulations underline the difficulty of reaching target 1 of the SDGs - ending extreme poverty. The simulated trends clearly show that the pace of poverty reduction declines as the poverty rate becomes low. In Latin America, a continuation of the present trend of household consumption growth will not be sufficient to lower the poverty rate below 1%, although poverty in 2011 is at 5% only. In the baseline scenario poverty declines to 1.89% within 19 years. At the same time, changes in inequality become much more important. Declining inequality would decrease the number of poor below 1% easily, while increasing inequality would see the number of extremely poor increase despite positive income growth. At some moment, growth is no longer sufficient when it comes to eliminating poverty, and additional policy measures will be necessary. Reaching the SDG target will be hardest in Sub-Saharan Africa. A continuation of the current positive trend is not going to be enough to bring poverty down to (close to) zero. Poverty will decline substantially, (by a little less than half in the baseline scenario) but then it is still at 24.9%. In East Asia poverty will be down to far less than 1%. Also in South Asia, poverty will decline substantially. A continuation of current trends will lower the poverty rate from currently 23% to below 4%. Nevertheless, there will still be many millions of people left in poverty.

In the last 20 years or so, there has been tremendous progress in reducing poverty. If we could keep up this rate of growth or even accelerate it, we can expect further good progress. Although it is unlikely that the first target of the SDGs will be reached everywhere, poverty will be greatly reduced. Also, as average incomes will hopefully grow, inequality and questions of distribution become more important for poverty reduction. Finally, eliminating extreme poverty everywhere is well within reach, although it might take a bit longer than until 2030.

References

Richard H. Adams. Economic growth, inequality and poverty: Estimating the growth elasticity of poverty. *World Development*, 32(12):1989–2014, 2004.

Abhijit Banerjee and Thomas Piketty. Top Indian Incomes, 1922-2000. *The World Bank Economic Review*, 19(1):1–20, 2005.

Timothy Besley and Robin Burgess. Halving global poverty. *The Journal of Economic Perspectives*, 17(3):3–22, 2004.

Surjit Bhalla. *Imagine There's No Country: Inequality, Poverty and Growth in the Era of Globalization.* Institute for International Economics, Washington DC, 2002.

Surjit S. Bhalla. Recounting the poor: Poverty in india, 1983-99. *Economic and Political Weekly*, 38(4):338–349, 2003.

Francois Bourguignon. The Growth Elasticity of Poverty Reduction: Explaining Heterogeneity across Countries and Time Periods. In T.S. Eischer and S.J. Turnovsky, editors, *Inequality and Growth: Theory and Policy Implications*, pages 3–26. MIT Press, Cambridge, MA and London, 2004.

Florent Bresson. On the Estimation of Growth and Inequality Elasticities of Poverty with Grouped Data. *Review of Income and Wealth*, June 2009.

Laurence Chandy, Natasha Ledlie, and Veronika Penciakova. The Final Countdown: Prospects for Ending Extreme Poverty by 2030. Policy paper, The Brookings Institution, 1775 Massachusetts Ave., Washington, DC, 04 2013.

Shaohua Chen and Martin Ravallion. How did the World's Poorest Fare in the 1990s? *Review of Income and Wealth*, 47(3):283–300, 2001.

Shaohua Chen and Martin Ravallion. The Developing World Is Poorer Than We Thought, But No Less Successful in the Fight against Poverty. *Policy Research Working Paper World Bank*, August 2008.

Angus Deaton. Adjusted Indian Poverty Estimates for 1999-2000. *Economic and Political Weekly*, pages 322–326, January 2003.

Angus Deaton. Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World). *The Review of Economics and Statistics*, February 2005.

Angus Deaton. *Measuring Poverty.* Oxford University Press, Oxford, 2006.

Angus Deaton and Valerie Kozel. Data and Dogma: The Great Indian Poverty Debate. *The World Bank research observer*, 20(2):177–199, 2005.

K. Deininger and L. Squire. New ways of looking at old issues: Inequality and growth. *Journal of Development Economics*, 57(2):259–287, 1998.

Fanny Dellinger and Giulia Diprossimo. The growth elasticity of poverty revisited. a comparison of three different estimation techniques. Seminar paper, Université Paris 1, Sorbonne-Panthéon, Paris, 2015.

Shatakshee Dhongde and Camelia Minoiu. Global poverty estimates: A sensitivity analysis. Working Paper 234, IMF International Monetary Fund, 2011.

Morten Jerven. The Relativity of Poverty and Income: How Reliable are African Economic Statistics? *African Affairs*, 109(434):77–96, 2009.

Nanak Kakwani. *Income Inequality and Poverty - Methods and Applications.* Oxford University Press, New York, 1980.

Nanak Kakwani. Poverty and Economic Growth - With Application to Cote d'Ivoire. *Living Standards Measurement Study*, Working Paper No. 63, February 1993.

Adriaan Kalwij and Arjan Verschoor. A Decomposition of Poverty Trends across Regions: The Role of Variation in the Income and Inequality Elasticities of Poverty, 2005.

Chakrangani Lenagala and Rati Ram. Growth Elasticity of Poverty: Estimates from New Data. *International Journal of Social Economics*, 37 (12):923–932, 2010.

Andy McKay. Growth and Poverty Reduction in Africa in the Last Two Decades: Evidence from an AERC Growth-Poverty Project and Beyond. *Journal of African Economies*, 22(suppl 1):i49–i76, 2013.

B. S. Minhas. Validation of Large Scale Sample Survey Data Case of NSS Estimates of Household Consumption Expenditure. *Sankhy: The Indian Journal of Statistics, Series B (1960-2002)*, 50(3):279–326, 1988.

Johan A. Mistiaen and Martin Ravallion. Survey Compliance and the Distribution of Income. Policy Research Working Paper WPS 2956, The World Bank Group, Washington, DC, 2003.

NSSO Expert Group. Suitability of Different Reference Periods for Measuring Household Consumption: Results of a Pilot Survey. *Economic and Political Weekly*, 38(4):307–321, 2003.

Rati Ram. Growth elasticity of poverty: Alternative estimates and a note of caution. *KYKLOS*, 57(4):601–610, 2006.

Martin Ravallion. Have We Already Met the Millennium Development Goal for Poverty? *Economic and Political Weekly*, 37(46):4638–4645, 2002.

Martin Ravallion. The Debate on Globalization, Poverty and Inequality: Why Measurement Matters. *International Affairs*, 79(4):739–753, 2003a.

Martin Ravallion. Measuring aggregate welfare in developing countries - How well do national accounts and surveys agree? *Review of Economics and Statistics*, 85(3):645–652, 2003b.

Martin Ravallion. How Long Will It Take to Lift One Billion People Out of Poverty? *Policy Research Working Paper World Bank*, January 2013.

Martin Ravallion and Shaohua Chen. What Can New Survey Data Tell Us about Recent Changes in Distribution and Poverty? *World Bank Economic Review*, 11(2):357–82, 1997.

Xavier Sala-i Martin. The world distribution of income: Falling poverty and...convergence, period. *The Quarterly Journal of Economics*, 121(2): 351–397, 2006.

Xavier Sala-i Martin and Maxim Pinkovskiy. African Poverty is Falling... Much Faster than You Think! Working Paper 15775, National Bureau of Economic Research, February 2010.

Joseph Stiglitz, Amartya Sen, and Jean-Paul Fitoussi. The Measurement of Economic Performance and Social Progress Revisited: Reflections and Overview. Technical report, OFCE - Centre de recherche en économie de Sciences Po, Quai d'Orsay 69, 75340 Paris, 2009.

K. Sundaram and Suresh D. Tendulkar. NAS-NSS Estimates of Private Consumption for Poverty Estimation: A Further Comparative Examination. *Economic and Political Weekly*, 38(4):376–384, 2003.

Miguel Székely and Marianne Hilgert. What's Behind the Inequality We Measure: An Investigation Using Latin American Data. Working Paper 409, Inter-American Development Bank, 1999.

UN-MDG Report. The Millennium Development Goals Report 2015. Technical report, United Nations Inter-Agency and Expert Group on MDG Indicators, 2015.

World Bank. *World Development Report 2000/2001 - Attacking Poverty.* The World Bank, 2000.

Yu Xie and Xiang Zhou. Income Inequality in Today's China. *PNAS*, 111 (19):6928–6933, 2014.

Appendix: Summaries

English Abstract

Abstract: I present a new and simple approach on how to forecast and replicate poverty trends based on the growth elasticity of poverty (GEP) and the inequality elasticity of poverty (IEP). I use an analytically derived elasticity and combine this with data from household surveys, i.e. mean income, inequality and growth rates are all taken from household survey data. The novelty of the approach is twofold: The income distribution is approximated by the Fisk distribution and instead of National Accounting System (NAS) data only household survey data is used. I perform an illustrative "test" of the model by replicating past poverty trends based on the survey data, and I find that the model generally performs well, (except for highly unequal middle income countries). I then apply the model to forecast future poverty trends and to assess whether the first target of the Sustainable Development Goals, i.e. to eliminate poverty until 2030, can be reached in the different world regions. My results show that in East Asia, the target will most likely be reached. Also in Latin America and South Asia, poverty is projected to be very low by 2030. It will be very difficult to reach the target in Sub-Saharan Africa.

German Abstract

Abstract: Eine neue, einfache Methode zum Prognostizieren und Replizieren von Armutstrends wird vorgestellt. Diese basiert auf der Wachstumselastizität der Armut und der Ungleichheitselastizität der Armut. Die Wachstumselastizität der Armut wird analytisch hergeleitet und mit Daten von Haushaltsbefragungen kombiniert. Dieser Ansatz ist aus zweierlei Gründen neu: Zur analytischen Herleitung der Elastizitäten wird die Einkommensverteilung durch die Fisk-Verteilung approximiert, und anstelle von Daten aus der VGR werden ausschließlich Haushaltsbefragungsdaten verwendet. Ein illustrativer "Test" der Methode wird durchgeführt, indem vergangene Armutstrends in 6 Ländern anhand der Haushaltsbefragungsdaten repliziert werden. Insgesamt funktioniert das gut – nur die Armutstrends von sehr ungleichen Ländern mittleren Einkommens bereiten Probleme. Die Methode wird dann dafür verwendet, Armutstrends bis 2030 für vier Weltregionen zu prognostizieren. Die Frage, ob das erste der Ziele der nachhaltigen Entwicklung erreicht werden kann, nämlich die extreme Armut bis 2030 zum Verschwinden zu bringen, wird untersucht. Die Arbeit zeigt auf, dass in Ostasien mit dem Erreichen dieses Ziels gerechnet werden kann, auch in Südasien und Lateinamerika ist das möglich, während in Sub-Sahara-Afrika das Beenden der Armut sehr schwierig wird.