Poverty Accounting

Richard Bluhm* Denis de Crombrugghe† Adam Szirmai‡

This version: March 9, 2018

Revised Manuscript

Abstract

This paper proposes a new framework for poverty accounting, that is, the decomposition of poverty into its proximate components. Using aggregated household surveys from 124 countries, we estimate the potential impacts of income growth and redistribution on poverty reduction, as well as their actual contributions to poverty reduction over the period from 1981 to 2010. Our fractional response approach shows that the potential impacts are highly non-linear and vary across regions and time. This non-linearity needs to be taken into account if empirical estimates are to inform development policy. Although historically growth has played the main role in poverty reduction, we find that initial inequality is a strong moderator of the impact of growth. In fact, there has been a shift towards pro-poor growth around the turn of the millennium, both at the $2 and at the $1.25 a day poverty line. Nevertheless, our projections of poverty rates until 2030 show that the end of extreme poverty within a generation, as put forth in the Sustainable Development Goals, is unlikely to materialize.

Keywords: poverty, inequality, growth, fractional response models
JEL Classification: I32, C25, O10, O15

This paper and previous versions have been presented at the World Bank, ECINEQ, Agence Française de Développement, Center for Global Development, IARIW, UGöttingen, OECD, UMainz, AFES, and several workshops in Paris, Leuven and Maastricht. We have greatly benefited from discussions with the participants. In particular, we would like to thank Florent Bresson, Laurence Chandy, Martin Gassebner, Gary Hoover, Mary Kaltenberg, Charles Kenny, Stephan Klasen, Melanie Krause, Nicolas Meisel, Pierre Mohnen, Christophe Muller, D.S. Prasada Rao, Thomas Roca, Susanne Steiner, Kaj Thomsson, Bart Verspagen and Jeffrey Wooldridge. We gratefully acknowledge financial support from the Agence Française de Développement (AFD). All remaining errors are unfortunately ours.

* *Corresponding author.* Leibniz University Hannover, Institute of Macroeconomics, and UNU-MERIT, Maastricht University; e-mail: *bluhm@mak.uni-hannover.de*

† Maastricht University; e-mail: *d.decrombrugghe@maastrichtuniversity.nl*

‡ UNU-MERIT, Maastricht University; e-mail: *szirmai@merit.unu.edu*

1 Introduction

How long will it take for economic growth to eradicate poverty? This question is at the heart of the ongoing debate about inclusive growth and equitable development. Okun's equity-efficiency trade-off, which for several decades seemed to override fairness concerns, has recently lost ground to a renewed focus on 'pro-poor growth' (World Bank, 2005) and 'shared prosperity' (World Bank, 2015). The shift in the policy discussion is supported by an increasingly large empirical literature analyzing the impact of changes in income and inequality on poverty, and their respective contributions towards poverty reduction.[1] Collectively, these studies not only confirm that income growth is crucial to sustained decreases in poverty, but also establish that the beneficial effect of growth strongly depends on the initial levels of income *and* inequality.[2]

In this paper we present a new framework for poverty accounting. Analogous with growth and development accounting, poverty accounting is the decomposition of (changes in) poverty into its proximate sources. It addresses questions such as, what is the impact of a one percent rise in average income or reduction in inequality on poverty rates? How much of the historical variation in poverty was due to economic growth, and how much to redistribution? While it is an established fact that changes in poverty rates are linked to changes in average incomes and the income distribution, there is no consensus on the appropriate functional form.

The cross-country literature usually approximates the relationship between poverty, average income and the income distribution using log-linear models with *ad hoc* interaction terms. The resulting fit is always moderate and the estimates are plausible for overall mean effects only. To policy makers, the mean effect is of limited interest since it is well known that, even within a country, growth can have a very different impact on poverty across regions or ethnic groups (Aaron, 1967, Hoover et al., 2008).

Our framework builds on the fractional response approach developed by Papke and Wooldridge (1996, 2008). The fact that the poverty headcount ratio is a fraction leads to a natural model of the expected poverty rate, $\mathrm{E}[H]$, given a fixed (absolute) poverty line, z. This model incorporates two crucial features: first, $\mathrm{E}[H]$ is bounded on the unit interval; and second, $\mathrm{E}[H]$ converges to zero (unity) if mean income becomes arbitrarily large (small) relative to the poverty line. It follows that elasticities and semi-elasticities of poverty with respect to income and inequality are highly non-linear. A central contribution of our paper is that we demonstrate the variety of analyses which can be consistently carried out using this framework. Specifically, we *(i)* estimate elasticities and semi-elasticities of poverty that remain plausible over the entire range of the data, *(ii)*

[1] See, for example, Ravallion and Chen (1997), Dollar and Kraay (2002), Besley and Burgess (2003), Kraay (2006), Kalwij and Verschoor (2007), and Dollar et al. (2016).

[2] In fact, this dependence arises mechanically, since poverty is functionally linked to average income and inequality (Datt and Ravallion, 1992, Kakwani, 1993, Bourguignon, 2003).

recover the conditional expectation function of the poverty headcount ratio, *(iii)* estimate the counterfactual quantities needed to compute the historical contributions of growth and redistribution to poverty reduction, and *(iv)* produce out-of-sample forecasts.

Estimating any poverty decomposition using summary household survey data entails a number of practical difficulties. A second contribution of this paper is that we extend the fractional response model to resolve three difficulties commonly encountered in applications: unobserved heterogeneity due to persistent measurement differences between surveys, unbalanced panel data due to infrequently undertaken surveys, and endogeneity due to time-varying measurement errors in per capita income or expenditure.

Our third and most substantive contribution is to apply the new framework to a large data set covering 124 countries over a 30-year period, in order to draw lessons for the future design of poverty reduction policies. In their most basic form, the data only contain three variables per country-survey-year: the poverty headcount ratio at a fixed international poverty line, average income or consumption expenditures, and the Gini coefficient. We demonstrate that our approach closely approximates the shape of the Lorenz curve near the poverty line using this limited information set.

Overall, we find that a one percent increase in mean income reduces the proportion of people living below the poverty line by about *two* percentage points, while a one percent increase in inequality raises the poverty rate by *one and a half* percentage points. We present estimates by region and time period, providing a basis for differentiated policy design. Regarding the historical contributions, we discover a shift in the poverty reducing pattern of growth around the turn of the millennium. Until 2000, about 90% of poverty reduction was due to income growth, and inequality tended to *rise* with higher growth rates. Since 2000, changes in inequality contributed almost a third to poverty reduction, and growth is no longer associated with rising inequality. We also present projections of the poverty headcount ratio until 2030 for two poverty lines: \$2 a day and \$1.25 a day (in 2005 PPPs). The \$2 a day poverty rate is likely to halve from about 40% in 2010 to below 20% in 2030, suggesting a billion people are being lifted out of poverty. At \$1.25 a day, however, the pace of poverty reduction is bound to slow down considerably.

Before we proceed, note that, just as growth accounting is silent on what ultimately drives growth, poverty accounting does *not* identify the ultimate determinants of poverty. Instead, it reveals how *given* changes in aggregate income and distribution translate into poverty outcomes. The causal effect of specific policies cannot be deduced from the decomposition of poverty into its proximate sources of change.

The remainder of the paper is organized as follows. Section 2 reviews how the existing literature decomposes poverty rates and estimates poverty elasticities. Section 3 introduces our approach and the econometrics of fractional response models. Section 4 briefly outlines the data used in this paper. Section 5 presents estimation results, elasticities, contributions, and poverty projections until 2030. Section 6 concludes.

2 Poverty decompositions and elasticities

We begin with a review of standard concepts and models used in the literature on poverty decompositions.

The poverty headcount ratio in some country at time t is defined as

$$H_t = H(\bar{y}_t/z, L_t) \equiv \Pr\{y_t \leq z\} \tag{1}$$

where $\bar{y}_t/z$ is the (relative) distance of mean income, $\bar{y}_t$, to the (fixed) poverty line, z; L_t is the Lorenz curve describing the entire income distribution; and y_t is the income of an individual drawn randomly from the population. Eq. (1) represents the probability that a random individual is poor, as a function of the mean level and distribution of income, as summarized by $\bar{y}_t$ and L_t. We may speak of a 'shortfall' of income when $\bar{y}_t < z$ and of 'affluence' when $\bar{y}_t > z$.

With micro-level data on incomes or consumption expenditures it is straightforward to decompose changes in H_t into changes in $\bar{y}_t$ and L_t (Datt and Ravallion, 1992, Kakwani, 1993).[3] Across countries, however, the underlying surveys are often not available but condensed in the form of tabulated (grouped) data. To overcome this limitation, Bourguignon (2003) suggests the adoption of a log-normal income distribution. This assumption is theoretically grounded and popular, yet it does have critics (e.g. Bresson, 2009) and will be relaxed later.

If individual income, y_t, is a log-normal random variable with $\ln y_t \sim \mathcal{N}(\mu_t, \sigma_t^2)$, then mean income can be written as $\bar{y}_t = \mathrm{E}[y_t] = \exp(\mu_t + \sigma_t^2/2)$, and the poverty headcount ratio at time t is

$$H_t = \Phi\left(\frac{-\ln(\bar{y}_t/z)}{\sigma_t} + \frac{1}{2}\sigma_t\right) \tag{2}$$

where $\Phi(\cdot)$ denotes the standard normal cdf. The income distribution is entirely determined by the standard deviation of log-income, σ_t. The derivatives of eq. (2) can be rearranged to yield the income and inequality elasticities of poverty, so that

$$\varepsilon_t^{H\bar{y}} \equiv \frac{\partial H_t}{\partial \bar{y}_t}\frac{\bar{y}_t}{H_t} = -\frac{1}{\sigma_t}\lambda\left(\frac{-\ln(\bar{y}_t/z)}{\sigma_t} + \frac{1}{2}\sigma_t\right) \tag{3}$$

and

$$\varepsilon_t^{H\sigma} \equiv \frac{\partial H_t}{\partial \sigma_t}\frac{\sigma_t}{H_t} = \left(\frac{\ln(\bar{y}_t/z)}{\sigma_t} + \frac{1}{2}\sigma_t\right)\lambda\left(\frac{-\ln(\bar{y}_t/z)}{\sigma_t} + \frac{1}{2}\sigma_t\right) \tag{4}$$

[3] A viable alternative to our method is to apply the micro-level approach to all countries by creating synthetic data as in Kraay (2006). It should be noted, though, that Lorenz curve estimation based on grouped data has its own problems (see Chotikapanich et al., 2007, Bresson, 2009, Krause, 2014).

where the inverse Mills ratio, $\lambda(x) \equiv \phi(x)/\Phi(x)$, is defined as the ratio of the standard normal pdf to the standard normal cdf, and we require $H_t > 0$. These analytic formulas inspired the econometric models of, among others, Bourguignon (2003), Kalwij and Verschoor (2007), and Klasen and Misselhorn (2008). To see this, note that the poverty rate can be decomposed in proportional terms using a Taylor linearization

$$\frac{\mathrm{d}H_t}{H_t} \approx \varepsilon_t^{H\bar{y}} \frac{\mathrm{d}\bar{y}_t}{\bar{y}_t} + \varepsilon_t^{H\sigma} \frac{\mathrm{d}\sigma_t}{\sigma_t} \tag{5}$$

where $\mathrm{d}H_t/H_t$, $\mathrm{d}\bar{y}_t/\bar{y}_t$ and $\mathrm{d}\sigma_t/\sigma_t$ are small relative changes in H_t, $\bar{y}_t$ and σ_t, respectively. Aitchison and Brown (1957) showed that, under log-normality, the standard deviation is a monotone transformation of the Gini coefficient of income inequality, G_t, namely

$$\sigma_t = \sqrt{2}\Phi^{-1}(G_t/2 + 1/2). \tag{6}$$

Eqs. (3) and (4) can therefore be used to estimate the elasticities using observed values of income and the Gini coefficient. With a little algebra, it is straightforward to derive an expression for the Gini elasticity and rewrite eq. (5) accordingly.

In empirical applications the poverty headcount equation is (usually) formulated in annualized differences, and the dependence of the elasticities on mean income and inequality is captured through interaction terms involving initial mean income and initial inequality. The result is what Bourguignon (2003) calls the 'improved standard model':

$$\begin{aligned}\Delta \ln H_{it} = \;& \alpha + \beta_1 \Delta \ln \bar{y}_{it} + \beta_2 \Delta \ln \bar{y}_{it} \times \ln(\bar{y}_{i,t-1}/z) + \beta_3 \Delta \ln \bar{y}_{it} \times \ln G_{i,t-1} \\ & + \gamma_1 \Delta \ln G_{it} + \gamma_2 \Delta \ln G_{it} \times \ln(\bar{y}_{i,t-1}/z) + \gamma_3 \Delta \ln G_{it} \times \ln G_{i,t-1} + \epsilon_{it}\end{aligned} \tag{7}$$

where Δ is the first-difference operator, α is a linear time trend, and ϵ_{it} is an error term; we also inserted the country subscript i.

The interaction terms in eq. (7) imply variable elasticities. The partial derivatives, $\varepsilon_{it}^{H\bar{y}} = \beta_1 + \beta_2 \ln(\bar{y}_{i,t-1}/z) + \beta_3 \ln G_{i,t-1}$ and $\varepsilon_{it}^{HG} = \gamma_1 + \gamma_2 \ln(\bar{y}_{i,t-1}/z) + \gamma_3 \ln G_{i,t-1}$, are sometimes referred to as the 'distribution-neutral' income elasticity and the 'growth-neutral' inequality elasticity. They serve as linear approximations to eqs. (3) and (4), offering simplicity but ignoring meaningful restrictions on the parameter space. In view of the bounded nature of the dependent variable ($H_{it} \in [0, 1]$), the functional form of the elasticities is predictably non-linear. Linear approximations must be poor for certain combinations of income and inequality away from the center. Yet to guide development policy we are precisely interested in regionally and temporally differentiated elasticities.

Eq. (7) has several other disadvantages. Differencing disregards the information contained in poverty *levels*, induces negative serial correlation, and is likely to compound measurement error.[4] The logarithmic transformation of H_t excludes poverty spells

[4] Consider a typical equation with well-behaved disturbances but a single serially correlated regressor

starting or ending with a zero, but discarding or modifying them affects consistency. The log transformation also makes it impossible to recover the conditional expectation of H_t except under strong independence assumptions, but the conditional expectation of H_t is what we need to compute counterfactuals and contributions to poverty reduction.[5] More subtly, it is unclear which level relationship might underlie eq. (7). Whereas differencing removes time-constant unobserved effects, the inserted interaction terms reintroduce the unobserved effects through the lagged levels. As a result, any systematic measurement difference between countries will bias the estimated coefficients.[6]

Last but not least, note that elasticities can paint a distorted picture of poverty dynamics. The income elasticity, for example, gives the impression that richer countries become ever better at poverty reduction because a drop in the poverty headcount from 2% to 1% is treated as equivalent to a drop from 50% to 25%. Hence, Klasen and Misselhorn (2008) recommend to focus on absolute changes in poverty rates. Removing the log from the headcount in eq. (7) achieves this. The partial effects become *semi*-elasticities, measuring *percentage point* changes in the number of poor expected for a given percentage change in income or inequality. Likewise, eqs. (3) and (4) provide semi-elasticities if we replace the inverse Mills ratio with the standard normal pdf. In contrast with elasticities, the semi-elasticities converge to zero as mean income becomes very large (or very small). For a policy maker who is concerned about the number of poor in the population, the semi-elasticities are the relevant parameters.

3 Fractional response models of poverty

3.1 Fractional probit

We now outline an approach to poverty accounting that does not suffer from the problems listed above. In a seminal paper, Papke and Wooldridge (1996) suggested modeling proportions using non-linear, parametric, fractional response models, known since as fractional logit, fractional probit etc. To the best of our knowledge, this idea has never been applied to poverty decompositions. The models are of the form $\mathrm{E}[y_i|\mathbf{x}_i] = F(\mathbf{x}_i'\boldsymbol{\beta})$, where $y_i \in [0,1]$ and $F(\cdot)$ is an invertible cdf, such as the standard normal. Thus we assume that the mean headcount ratio can be described by

$$\mathrm{E}[H_{it}|\bar{y}_{it}, G_{it}] = \Phi(\alpha + \beta \ln \bar{y}_{it} + \gamma \ln G_{it}) \quad \text{for} \quad i = 1, \ldots, N;\ t = 1, \ldots, T \tag{8}$$

affected by serially uncorrelated measurement error. It can be verified that first-differencing not only induces negative serial correlation in the disturbance term, but also magnifies attenuation bias.

[5]Moreover, as Santos Silva and Tenreyro (2006) point out for the case of gravity models, heteroskedasticity in the level equation causes bias in the parameter estimates after taking logs.

[6]As a case in point, Kalwij and Verschoor (2007) present a model in differences with unobserved effects removed, but later estimate interaction models with unobserved effects again present though not accounted for.

where $\ln z$ is absorbed into the constant and both explanatory variables are expressed in logs to linearize the index and ease interpretation.[7] Naturally, we expect $\beta < 0$ and $\gamma > 0$. We will temporarily ignore econometric complications such as unobserved heterogeneity, unbalanced data and endogeneity, to be discussed later.

[Figure 1 about here]

Eq. (8) may be rewritten as $\Phi^{-1}\left(\mathrm{E}[H_{it}|\bar{y}_{it}, G_{it}]\right) = \alpha + \beta \ln \bar{y}_{it} + \gamma \ln G_{it}$, where $\Phi^{-1}(\cdot)$ plays the role of a 'link function' in the spirit of the Generalized Linear Model (GLM) literature. Figure 1 plots $\Phi^{-1}(H_{it})$ against $\ln \bar{y}_{it}$ per region, including a regression line. It is striking how closely the inverse normal cdf linearizes the relationship.[8] Rather than transforming the data, though, it is preferable to estimate the conditional expectation function given in eq. (8) directly. A suitable method is Quasi Maximum Likelihood (QML) using a Bernoulli likelihood. QML is as robust as non-linear least squares, but potentially more efficient. Most importantly, it only requires correct specification of the conditional mean, irrespective of the true distribution of H_{it} (Gourieroux et al., 1984).

Plugging in the QML estimates, the income and Gini elasticities are as follows

$$\hat{\varepsilon}_{it}^{H\bar{y}} = \frac{\partial \hat{\mathrm{E}}[H_{it}|\bar{y}_{it}, G_{it}]}{\partial \bar{y}_{it}} \times \frac{\bar{y}_{it}}{\hat{\mathrm{E}}[H_{it}|\bar{y}_{it}, G_{it}]} = \hat{\beta} \times \lambda\left(\hat{\alpha} + \hat{\beta} \ln \bar{y}_{it} + \hat{\gamma} \ln G_{it}\right) \tag{9}$$

and

$$\hat{\varepsilon}_{it}^{HG} = \frac{\partial \hat{\mathrm{E}}[H_{it}|\bar{y}_{it}, G_{it}]}{\partial G_{it}} \times \frac{G_{it}}{\hat{\mathrm{E}}[H_{it}|\bar{y}_{it}, G_{it}]} = \hat{\gamma} \times \lambda\left(\hat{\alpha} + \hat{\beta} \ln \bar{y}_{it} + \hat{\gamma} \ln G_{it}\right). \tag{10}$$

Eqs. (9) and (10) closely mimic the properties of the analytical elasticities derived from the log-normality assumption in eqs. (3) and (4). They will approach zero whenever the inverse Mills ratio approaches zero, and remain plausible for extreme values of the covariates. The non-linearity arises directly from the bounded nature of the headcount ratio. This secures a number of advantages: the information contained in poverty levels is not wasted, the model will predict poverty headcount ratios strictly within the unit interval, and elasticities as well as semi-elasticities are estimated consistently. By plugging counterfactual values for $\bar{y}_{it}$ and G_{it} into $\hat{\mathrm{E}}[H_{it}|\bar{y}_{it}, G_{it}]$, we can estimate the respective contributions of growth and redistribution to poverty reduction. Similarly, given growth and inequality scenarios, we can predict the future path of poverty rates.

Note that we *do not require log-normality.* Any two-parameter distribution will do, as long as the poverty headcount remains a smooth function of mean income and inequality,

[7] Simulations (analogous to those presented in Table C-2 of Online Appendix C) suggest that using logs greatly improves the ability of the model to accurately recover known (semi-)elasticities.

[8] The interactions used in the standard log-linear models do not fully capture this intrinsic non-linearity. Figure A-2 in Online Appendix A shows the same graph using the log of the Gini coefficient.

up to statistical error. In fact, it is possible to nest the log-normal case within the GLM.[9]

3.2 Heterogeneity, unbalanced data, and endogeneity

We now discuss the three technical problems we encounter in our application: unobserved heterogeneity, unbalanced data, and measurement error. For the purpose and duration of this subsection, we stack the *time-varying* covariates and their coefficients in the vectors $\mathbf{x}_{it} = (x_{it,1}, \ldots, x_{it,K})'$ and $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_K)'$, and form the matrices $\mathbf{X}_i = (\mathbf{x}_{i1}, \ldots, \mathbf{x}_{iT})'$. The ideal model we would like to estimate is

$$\mathrm{E}[H_{it}|\mathbf{X}_i, \eta_i] = \mathrm{E}[H_{it}|\mathbf{x}_{it}, \eta_i] = \ \Phi(\mathbf{x}_{it}'\boldsymbol{\beta} + \eta_i) \quad \text{for} \quad i = 1, \ldots, N; \ t = 1, \ldots, T \tag{11}$$

where the covariates are strictly exogenous conditionally on unobserved country-level effects η_i, and the panel is balanced. The unobserved effects are meant to capture time-persistent differences in measurement or deviations from a two-parameter distribution, which may be arbitrarily correlated with elements of $\mathbf{X}_i$. The problem, in our small T, large N context, is that the η_i are incidental parameters and cannot be consistently estimated. The estimation of $\boldsymbol{\beta}$ as well as all partial effects is affected.[10]

As a solution, Papke and Wooldridge (2008) propose a *correlated random effects* (CRE) structure, using the Mundlak (1978) and Chamberlain (1984) device: $\eta_i|(\mathbf{x}_{i1}, \ldots, \mathbf{x}_{iT}) \sim \mathcal{N}(\varphi + \bar{\mathbf{x}}_i'\boldsymbol{\theta}, \sigma_u^2)$ where $\bar{\mathbf{x}}_i = T^{-1}\sum_{t=1}^{T}\mathbf{x}_{it}$ contains time averages of *all* time-varying regressors $\mathbf{x}_{it}$. Defining $u_i \equiv \eta_i - \varphi - \bar{\mathbf{x}}_i'\boldsymbol{\theta}$, it follows that $u_i|(\mathbf{x}_{i1}, \ldots, \mathbf{x}_{iT}) \sim \mathcal{N}(0, \sigma_u^2)$. Plugging this into eq. (11) and marginalizing with respect to u_i, we obtain what is known in the GLM literature as a 'population-averaged model'

$$\mathrm{E}[H_{it}|\mathbf{X}_i] = \mathrm{E}[\Phi(\varphi + \mathbf{x}_{it}'\boldsymbol{\beta} + \bar{\mathbf{x}}_i'\boldsymbol{\theta} + u_i)|\mathbf{X}_i] = \Phi\left(\varphi_u + \mathbf{x}_{it}'\boldsymbol{\beta}_u + \bar{\mathbf{x}}_i'\boldsymbol{\theta}_u\right). \tag{12}$$

The subscript u indicates that the coefficients have been rescaled by the factor $(1+\sigma_u^2)^{-1/2}$. The scaled coefficients and average partial effects (APEs) of all time-varying covariates are identified and can be consistently estimated.[11] Note that the linear counterpart of this specification would deliver estimates which are numerically identical to those of a standard fixed-effects model.

A second difficulty is to account for unbalanced data.[12] Let T_i denote the sample size for country i, with $T = \max_i T_i$, and let the selection of observations, as recorded by selection indicators s_{it}, be ignorable (conditionally independent). Unbalancedness implies

[9] A dedicated section of Online Appendix B discusses this aspect in more detail.

[10] The biases decrease as T gets larger, but there are no benchmark simulations for the fractional probit case that we know of, and in our application the sample sizes are definitely small.

[11] As the method is QML, asymptotic standard errors based on the inverse information matrix will be too conservative (see for details Papke and Wooldridge, 2008, and Wooldridge, 2010a). Standard errors should be adjusted and made robust to clustering at the country level.

[12] Unbalancedness is a serious issue in our application, see Table A-1 in Online Appendix A.

that the heterogeneity depends on T_i. Wooldridge (2010a) argues that this dependence can be captured by allowing both the mean and the variance of the unobserved country effects, η_i, to differ per sample size. Hence we define T dummy variables $\delta_{T_i,\ell}$ using the Kronecker delta ($\delta_{T_i,\ell} = 1$ if $T_i = \ell$ and 0 otherwise). We include these dummies as well as interactions with the time averages ($\delta_{T_i,\ell}\bar{\mathbf{x}}_i'$) among the covariates in eq. (12), and scale the argument of $\Phi(\cdot)$ by the square root of $\mathrm{Var}(\eta_i) = \exp\left(2\sum_{\ell=2}^{T-1}\delta_{T_i,\ell}\omega_\ell\right)$, where the ω_ℓ are unknown variance parameters. The result is the following heteroskedastic model

$$\mathrm{E}[H_{it}|\{s_{it}, s_{it}\mathbf{x}_{it}; t = 1, ..., T\}] \;=\; \Phi\left(\frac{\mathbf{x}_{it}'\boldsymbol{\beta}_h + \sum_{\ell=2}^{T}\delta_{T_i,\ell}\left(\varphi_{h\ell} + \bar{\mathbf{x}}_i'\boldsymbol{\theta}_{h\ell}\right)}{\exp\left(\sum_{\ell=2}^{T-1}\delta_{T_i,\ell}\omega_\ell\right)}\right) \tag{13}$$

where the subscript h denotes a new scale factor and $T_i \geq 2$. This specification nests the balanced case. In that case, the numerator contains only one set of time averages and a constant in addition to the time-varying covariates, and the denominator is unity.

Endogeneity due to measurement error in income or consumption expenditures adds another layer of complexity. Ravallion (2003) and Deaton (2005), among others, discuss the presence of systematic biases in survey means and variances. If such biases are random (of the classical measurement error type), they attenuate the income coefficient, but if they are survey-specific (non-classical), they may work in the opposite direction (Ravallion and Chen, 1997). When income and poverty are estimated from the same survey, then errors in the former entail errors in the latter. How inference is affected depends on which type of error is stronger and how this spills over into other variables.[13] Hence, like Ravallion (2001) and Kalwij and Verschoor (2007), we treat the observed mean income ($\ln \bar{y}_{it}$) as endogenous.

We cannot simply replace observed mean incomes by a projection in non-linear models. Instead, we follow Heckman and Robb (1985) who describe an alternative two-step 'control function' method analogous to the standard instrumental variable method in linear models. Intuitively, rather than using the prediction, the reduced-form residual of an endogenous regressor is included (as a 'control' variable) to capture the endogenous variation in the equation of interest. Such a method is developed by Papke and Wooldridge (2008) for balanced panels and can be extended to the unbalanced case. We require $m \geq 1$ time-varying external instruments, relevant but strictly exogenous conditionally on the unobserved effects, which we arrange in vectors $\mathbf{z}_{it}$, with time averages $\bar{\mathbf{z}}_i$, and matrices $\mathbf{Z}_i = (\mathbf{z}_{i1}, \ldots, \mathbf{z}_{iT})'$. The first step is to estimate a reduced form $\ln \bar{y}_{it} = \mathbf{x}_{1it}'\boldsymbol{\pi}_1 + \mathbf{z}_{it}'\boldsymbol{\pi}_2 + \sum_{\ell=2}^{T}\delta_{T_i,\ell}(\pi_{0\ell} + \bar{\mathbf{x}}_i'\boldsymbol{\pi}_{3\ell} + \bar{\mathbf{z}}_i'\boldsymbol{\pi}_{4\ell}) + \nu_{it}$, where $\mathbf{x}_{1it}$ is $\mathbf{x}_{it}$ *excluding* the mismeasured $\ln \bar{y}_{it}$ but *including* time dummies, and $\bar{\mathbf{x}}_i$ includes the corresponding time

[13] Classical and non-classical measurement error in income may affect the Gini too (Chesher and Schluter, 2002). We reach no simple conclusion about the direction of biases, but note that the correlation between the Gini coefficient and average income in our sample is practically zero.

averages. The second step is to estimate the population-averaged model conditioned on the exogenous variables *and* the reduced-form residuals, $\hat{\nu}_{it}$:

$$\mathrm{E}[H_{it}|\mathcal{S}] = \Phi\left(\frac{\mathbf{x}'_{1it}\boldsymbol{\beta}_g + \psi_g \ln \bar{y}_{it} + \rho_g \hat{\nu}_{it} + \sum_{\ell=2}^{T} \delta_{T_i,\ell}\left(\varphi_{g\ell} + \bar{\mathbf{x}}'_i \boldsymbol{\theta}_{g\ell} + \bar{\mathbf{z}}'_i \boldsymbol{\zeta}_{g\ell}\right)}{\exp\left(\sum_{\ell=2}^{T-1} \delta_{T_i,\ell}\omega_\ell\right)}\right) \tag{14}$$

where $\mathcal{S} = \{s_{it}, s_{it}\mathbf{x}_{it}, s_{it}\mathbf{z}_{it}; t = 1, ..., T\}$, the subscript g denotes another scale factor, and $T_i \geq 2$. Both steps use the Mundlak-Chamberlain device for all strictly exogenous variables. Asymptotic standard errors can be approximated via a panel bootstrap.[14] To assess the endogeneity of $\ln \bar{y}_{it}$ in eq. (14) we test $\rho_g = 0$, which amounts to a Hausman test of the augmented-regression type (see Hausman, 1978, Lin and Wooldridge, 2017).[15]

3.3 Elasticities and semi-elasticities

By plugging QML estimates into eq. (14) and averaging over the cross-sectional dimension, we obtain what Blundell and Powell (2004) and Wooldridge (2010a,b) call the *Average Structural Function* (ASF).[16] The derivative of the ASF with respect to a continuous variable is the *Average Partial Effect* (APE) of that variable. These APEs rather than the model parameters are the quantities of interest. In our model, where $\bar{y}_{it}$ and G_{it} appear in logs and do not show up in the variance function, the APEs correspond to *semi*-elasticities. More formally, denote the linear predictors inside the cdf in eq. (14) by $\mathbf{m}'_{it1}\hat{\boldsymbol{\xi}}_1$ for the numerator and $\mathbf{m}'_{it2}\hat{\boldsymbol{\xi}}_2$ for the denominator, then the estimated APE or semi-elasticity for variable x_k at time t is

$$\mathrm{APE}_t(x_k) = \hat{\eta}_t^{Hx_k} = \hat{\xi}_{1k} \times N^{-1} \sum_{i=1}^{N} \exp\left(-\mathbf{m}'_{it2}\hat{\boldsymbol{\xi}}_2\right) \phi\left(\mathbf{m}'_{it1}\hat{\boldsymbol{\xi}}_1 / \exp(\mathbf{m}'_{it2}\hat{\boldsymbol{\xi}}_2)\right). \tag{15}$$

The elasticity ($\hat{\varepsilon}_t^{Hx_k}$) is obtained by replacing $\phi(\cdot)$ by $\lambda(\cdot)$ in eq. (15).

We conduct several Monte Carlo experiments studying the magnitude of the errors in the estimated (semi-)elasticities and predicted poverty rates. Simulating data from various known distributions, we find small approximation errors for the QML approach,

[14] We implement the method in a Stata module called `fhetprob` with analytic first and second derivatives, see `www.richard-bluhm.com/data/`. An interesting extension is to let the conditional variance depend on inequality, implying a marginal proportional rate of substitution ($-\varepsilon_t^{H\bar{y}}/\varepsilon_t^{HG}$) that is variable rather than constant.

[15] The test requires controlling for country-specific heterogeneity in $\hat{\nu}_{it}$ in order to separate endogeneity due to correlation with the unobserved heterogeneity from endogeneity due to correlation with the time-varying idiosyncratic error (Lin and Wooldridge, 2017). The term $\bar{\mathbf{x}}'_i$ in eq. (14) achieves this, since it contains the individual time average of $\ln \bar{y}_{it}$.

[16] The ASF deals with the unobserved effects, varying panel sizes, and endogeneity by averaging eq. (14) over the cross-sectional dimension. Since the surveys are irregularly spaced and often cover only parts of a given year, the time path of the ASF is hard to track. Hence we also average over time, and obtain a single ASF for each variable of interest.

irrespective of whether we focus on (semi-)elasticities or predicted poverty rates. The traditional linear models make large errors with biases in the estimated (semi-)elasticities exceeding 100% at very low or very high incomes. The results of these experiments are summarized in Online Appendix C.

4 Data

We compile a data set covering 124 countries over 30 years (1981-2010), based on household surveys from the World Bank's *PovcalNet* database.[17] Subsamples of this data have been used in previous studies (e.g. Adams, 2004, Kalwij and Verschoor, 2007, Chambers and Dhongde, 2011) and the World Bank's methodology is described in detail in Chen and Ravallion (2010). Here we will only mention its main features.

The information provided by *PovcalNet* consists of poverty headcount ratios (H_{it}), per capita monthly income or consumption expenditures ($\bar{y}_{it}$), the Gini coefficient of inequality (G_{it}), and the total population (pop_{it}). We consider two different poverty lines (z) widely used for international comparisons: z = \$2 a day (\$60.83 a month) and z = \$1.25 a day (\$38 a month). The latter is typically used to assess *extreme* poverty. In addition, we define dummy variables indicating whether $\bar{y}_{it}$ denotes income or consumption, and whether it is reported at the level of individual households or has been tabulated for groups of households (deciles or finer quantiles). Reported poverty rates tend to be lower in income surveys than in expenditure surveys, and the availability of individual versus grouped data may explain other systematic differences. About 63% of the data come from expenditure surveys and about 74% are estimated from grouped data. All amounts are in constant international dollars at 2005 PPP-adjusted prices.

The surveys in the *PovcalNet* database are usually designed to be nationally representative, but three large countries (China, India and Indonesia) keep urban and rural surveys separate. To construct national series in those cases we weight the poverty and income data using urban and rural population shares. Since the Gini coefficient is not subgroup-decomposable, we estimate a national Gini coefficient via an approximation due to Young (2011), based on a mixture of two log-normal distributions.[18] If in a given year only an urban or rural survey is available we drop it, except in the case of Argentina where urbanization is close to 90% and we treat the urban series as nationally representative. We end up with an unbalanced panel of 124 countries spanning 30 years, with an average

[17]The data is publicly available at `http://iresearch.worldbank.org/PovcalNet` (we used the 2005 PPP version). Per capita averages are simple averages without equivalence scaling.

[18]*PovcalNet* omits weighting some recent data. For the sake of consistency, we apply this approximation in all cases where separate urban and rural surveys must be combined. The formula is $G = \sum_{i=1}^{K}\sum_{j=1}^{K}\frac{w_i w_j \bar{y}_i}{\bar{Y}}\left(2K\left[\frac{\ln\bar{y}_i - \ln\bar{y}_j + \frac{1}{2}\sigma_i^2 + \frac{1}{2}\sigma_j^2}{(\sigma_i^2+\sigma_j^2)^{1/2}}\right] - 1\right)$, where K is the number of subgroups, w_i is the population share of the i-th subgroup, $\bar{y}_i$ is its mean income, σ_i^2 is the corresponding income variance, and $\bar{Y}$ is the population-weighted mean income across subgroups.

time series length ($\bar{T}$) of about 6.5 surveys and a total of 809 observations. Table 1 provides summary statistics for the entire panel.[19]

[Table 1 about here]

The linear approximations used in the extant literature require data in differences or log-differences. For the sake of comparison, we calculate annualized survey-to-survey differences. Some of these differences coincide with switches between income and expenditure surveys or changes in the level of aggregation; we exclude those and only retain differences that are both income or expenditure based, and both derived from micro data or tabulated data. Since differencing requires $T_i \geq 2$ and we only use comparable surveys, we end up with 648 differences from 104 countries. To estimate historical contributions we reduce this data set further, keeping only the longest consecutive runs of comparable surveys (yielding 123 observations). We later also consider two sub-periods, preceding and following the year 2000 (yielding 87 observations in each).

Apart from the survey-based data, we collect consumption data from national accounts, which will serve both as an instrument for survey income and as a basis for the projections. Per capita personal consumption expenditures (PCE) are retrieved from both the World Development Indicators (WDI) and the Penn World Table 7.1 (PWT) and combined into a single series.[20] Population projections for the period 2010-2030 are taken from the World Bank's Health, Nutrition and Population Statistics database.

5 Results

5.1 Regressions

Table 2 presents our main estimation results, with each specification addressing an additional issue: unobserved effects, unbalancedness and measurement error, in that order. All regressions include time averages à la Mundlak to capture measurement differences across countries (unobserved effects), and survey type dummies (expenditures or income, grouped or household data) to capture measurement differences across surveys. In addition, a full set of year dummies allows for unspecified common time effects.

[Table 2 about here]

[19] Online Appendix A provides further summary statistics. Table A-1 illustrates the survey coverage over time and indicates the different survey types. Table A-2 reports summary statistics by region.

[20] Monthly PCE^P_{it} is computed as $(kc_{it}/100 \times rgdpl_{it}/12)$, where kc_{it} is the PWT consumption share and $rgdpl_{it}$ is GDP per capita (Laspeyres) in 2005 constant prices. PCE^W_{it} is WDI household final consumption expenditures in 2005 prices divided by population and converted to monthly figures. The 'merged' series is constructed using the WDI as the default but replacing it with PWT data if coverage over 1981-2010 is better.

Column (1) includes correlated random effects, but ignores unbalancedness; see eq. (12). The coefficients have the expected signs but an arbitrary scale, which is why the adjacent column reports average partial effects (APEs) averaged over time. These APEs should be interpreted as *semi-elasticities*. If mean income grows by one *percent* we expect the poverty rate to go down by 0.284 *percentage points*, and if the Gini increases by one *percent* we expect the poverty rate to go up by 0.232 *percentage points*. The corresponding average *elasticities* have higher standard errors (se), and are about -1.83 ($se = 0.084$) for average income and 1.50 ($se = 0.167$) for the Gini. Note that these values are close to the lower end (in absolute value) of the range reported in the literature.[21]

This first specification could be affected by the severe unbalancedness of the panel and the presence of measurement error in income. Column (2) addresses the unbalancedness by including panel size dummies, interactions between the Mundlak time averages and the panel size dummies, and a variance function (in the denominator, see eq. (13)). The substantive conclusions hardly change. The APE of income remains almost identical and the APE of inequality increases by less than one standard error. Thus, unbalancedness seems to cause little bias on average, though it may still affect the (semi-)elasticities at particular points in time and space.

Our preferred specification, column (3), is the empirical counterpart of the two-step estimator presented in eq. (14). To account for measurement error (hence endogeneity) in income, we instrument survey mean income or expenditure with per capita consumption expenditures from the PWT national accounts (PCE_{it}^P). The main identifying assumption is that any measurement error in per capita consumption from the national accounts is uncorrelated with survey-based measurement error in mean income or expenditures (see also Ravallion, 2003, Deaton, 2005).[22]

The evidence favors endogeneity in income (panel bootstrap t-stat ≈ 1.79), especially if we ignore the first-stage sampling error which plays no role under the null hypothesis of exogeneity (t-stat ≈ 2.04). The APE and income elasticity are now larger in absolute value than in the previous two specifications ($\bar{\varepsilon}^{H\bar{y}} \approx -2.24$, $se = 0.244$). A one percent increase in average income raises the expected reduction in the poverty rate by 0.338 *percentage points* or 2.24 *percent*. We tentatively conclude that the coefficient of income in columns (1) and (2) is attenuated. Attenuation seems to be more of a problem than systematic survey bias, although we cannot rule out that more complex error structures are at play. As expected, the APE and elasticity of inequality are practically unaffected by this correction ($\bar{\varepsilon}^{HG} \approx 1.70$, $se = 0.253$).[23]

[21] The typical range for the income elasticity in earlier studies is from about -2 to -5, that for the Gini elasticity is even wider. More recent studies report income elasticities closer to -2.

[22] Figure B-1 in Online Appendix B shows a partial regression plot highlighting the strength of the first stage relationship.

[23] We recognize that inequality as recorded in the Gini may be affected by measurement error too (see footnote 13), but we lack an adequate instrument for the Gini coefficient.

The last row of Table 2 reports the square root of the mean squared residual ($\sqrt{MSE}$). The poverty headcount is predicted with about three and a half percentage points accuracy in the first model, and with better than two and a half percentage points accuracy in the next two. A close fit is what one would expect from a well-defined decomposition. A simple pseudo-R^2 measure, the squared correlation between the observed and fitted values, with values of 0.98 to 0.99, tells the same story.

We compare the results presented here with those obtained by applying the traditional linear approaches to our data in Online Appendix B. To summarize, the traditional linear models perform poorly and are unlikely to produce reliable estimates over a wide range of circumstances. We also examine the evidence in favor of log-normality in our data, and find that this hypothesis is consistently rejected.

5.2 Impacts

The strength of the fractional response approach lies in its ability to deliver precise and unbiased estimates of effects other than the overall mean response. Table 3 and Table 4 illustrate this point by estimating income and Gini elasticities, Panel (a), and semi-elasticities, Panel (b), over different time periods for six large geographic regions. Semi-elasticities are computed according to eq. (15), by plugging in time-period and region-specific averages of mean income ($\ln \bar{y}_{it}$) and inequality ($\ln G_{it}$), and then averaging over the entire subsample. The calculation of the elasticities is analogous, with $\lambda(\cdot)$ replacing $\phi(\cdot)$ in eq. (15). Standard errors are computed via a panel bootstrap to take into account the sampling uncertainty of the first stage.

[Table 3 about here]

There is considerable regional and temporal variation in the estimated income elasticities. As argued in Section 2, heterogeneity of elasticities is a consequence of heterogeneity in income and inequality. More affluent regions (Eastern Europe and Central Asia, Latin America and the Caribbean, and the Middle East and North Africa) have higher income elasticities (in absolute value) than poorer regions (East Asia and Pacific, South Asia, and Sub-Saharan Africa). Income dynamics over time are also clearly visible. In Eastern Europe and Central Asia, for example, income is comparatively high before the post-communist transition, sharply collapses through the 1990s, and recovers during the 2000s. The income elasticity follows suit. Compared to earlier results (e.g. Kalwij and Verschoor, 2007), we find markedly higher average income elasticities in more affluent regions and lower elasticities in poorer regions. They are also estimated with much greater precision. All standard errors in Table 3 are small compared to the point estimates and *remain* small for regions with extreme values.

The picture is reversed when we consider the semi-elasticities in Panel (b). Comparatively affluent regions have fewer people near the poverty line, and thus the poverty reduction potential from an equivalent increase in average income is much smaller in terms of people lifted out of poverty. This pattern is (again) best visible in Eastern Europe and Central Asia, where absolute poverty at the $2 a day poverty line is almost non-existent just before the post-communist transition, but rises sharply in the 1990s as incomes decline. Correspondingly, the semi-elasticity is close to zero in the 1980s but then increases as more people fall into poverty. Likewise, the biggest poverty reduction potential in 2005-2010 was in East Asia, South Asia, and Sub-Saharan Africa. This highlights an important point. For development policy, we really care about the share of the population lifted out of poverty rather than the percent change in the poverty rate.

[Table 4 about here]

Turning to the Gini elasticity, Panel (a) of Table 4 shows where the potential of redistributive policies in terms of relative reductions of the poverty rate has been largest in the last three decades. Unsurprisingly, these regions are Eastern Europe and Central Asia, Latin America and the Caribbean, and the Middle East and North Africa – all of which have above average inequality. Sub-Saharan Africa starts out with high inequality (the population-weighted mean Gini in the 1980s is 0.46) and very low income, so that the Gini elasticity is small. This is the downside of the dependency on initial levels: countries can be so poor and unequal that the immediate effects of equalization and income growth on *relative* changes in the poverty rate are small. Again, though, the semi-elasticities presented in Panel (b) reverse the picture: poorer and richer countries swap positions. The potential for *absolute* reductions in poverty rates through redistribution was larger in poorer regions throughout the period.[24]

Two features stand out in these regional and temporal comparisons. First, there is the contrast between the two impact measures: elasticities and semi-elasticities. The latter have the higher policy relevance. Second, there is the sizable heterogeneity in the estimated impacts across space and time, mainly due to their dependence on the prevailing levels of income and inequality.[25] Figure 2 illustrates this by graphing the estimated poverty elasticities and semi-elasticities. The functions are computed from eq. (15) and its analogue with $\lambda(\cdot)$ replacing $\phi(\cdot)$, by plugging in various combinations of values for per capita income ($\ln \bar{y}_{it}$) and inequality ($\ln G_{it}$), and then averaging over the entire sample. As Figure 2a shows, on top of the direct poverty alleviating effect of income redistribution, a lower level of inequality also raises the income elasticity in absolute value at every point. The magnitude of both elasticities is steeply increasing in the level of income. The returns to either income growth or equalization are bigger, the higher the

[24] A dedicated section in Online Appendix B shows that all of these results also hold at the $1.25 line.

[25] Incomes (or expenditures) have increased substantially in all regions between 1981 and 2010 (see Table A-3 in Online Appendix A). By contrast, inequality shows no systematic trend over the sample period from 1981 to 2010. In a simple regression of the Gini coefficient on time, we fail to reject the null hypothesis that the time trend is zero (cluster-robust t-stat ≈ 0.07 and $p > 0.94$).

income level, and the gap between the functions evaluated at different inequality levels keeps widening, inviting the conclusion that redistribution has most impact in affluent societies. That, precisely, is the misleading feature of poverty elasticities.

[Figure 2 about here]

Figure 2b shows the predicted income and Gini *semi*-elasticities of poverty. The picture is very different and in many ways more intuitive. If the income shortfall is too large – the mass of the income distribution is far below the poverty line – then both the income and the Gini semi-elasticities approach zero. If the country is affluent – the mass of the income distribution is far above the poverty line – then both semi-elasticities again approach zero. In between those two extremes, improvements in the income distribution can make a large difference in the proportion of people lifted out of poverty, both directly through redistribution and indirectly through growth (see also Bourguignon, 2003, Kalwij and Verschoor, 2007). This pattern leads us to the conclusion that poverty reduction strategies should focus both on income growth *and* on equalization, especially in low-income countries where the total returns to redistribution are large.[26]

An open question is whether this decomposition can be further improved by including additional, more 'ultimate' determinants of poverty, such as institutions, human capital, access to credit, or trade openness. Weak institutions, for example, may help explain why a disproportionate share of the gains of growth go to the rich. Under log-normality mean income and the Gini fully describe the distribution of incomes or expenditures, but we deliberately do not rely on log-normality and our data reject it. More realistic distributions use additional shape parameters in order to better capture skewness, a long tail, or multiple modes. In Online Appendix B we show that the inclusion of additional covariates has virtually no effect on the estimated decomposition. We conclude that the specification with only income and the Gini coefficient, plus panel size dummies and correlated random effects, is essentially saturated.

5.3 Contributions

The contribution of growth or redistribution to changes in poverty over some predefined period is the product of their impact (partial effect) and their actual variation, usually expressed as a proportion of the total change in the poverty rate over the same period. Authors performing such calculations tend to highlight the primacy of growth (see e.g. Dollar and Kraay, 2002, Kraay, 2006, Kraay et al., 2014, Dollar et al., 2016). Kraay (2006), for example, argues that long-run differences in poverty reduction can largely be explained by growth in average incomes, whereas changes in inequality seem to matter

[26] This finding should stand with the new 2011 PPPs as well. Different PPPs would shift incomes, the poverty line and the poverty rates, and so countries would be located differently on the graph; but the position of the curves in terms of relative income, $\bar{y}_t/z$, would be comparable.

only in the short run. Indeed, the greater importance of growth is largely due to the stylized facts that *(i)* there is a lot more long-run variation in growth than in inequality, and *(ii)* growth tends to be distribution-neutral.[27]

In spite of these facts, the poverty reducing quality of growth is of legitimate concern. While the cross-country correlation between growth and changes in inequality is near zero on average, it varies substantially across countries, regions, and time periods (Ravallion and Chen, 2003). In fact, the variability of the (semi-)elasticities emphasized above, combined with the variability of average income and inequality, can give rise to all sorts of relative contributions in individual countries or regions at particular points in time. This is the type of variation we are interested in.

Some additional notation will prove useful. Following Datt and Ravallion (1992), a discrete-time, non-logarithmic version of the decomposition of poverty changes is

$$\Delta H_t \approx \Big[H(\bar{y}_t/z, L_{t-1}) - H(\bar{y}_{t-1}/z, L_{t-1})\Big] + \Big[H(\bar{y}_{t-1}/z, L_t) - H(\bar{y}_{t-1}/z, L_{t-1})\Big]. \qquad (16)$$

Here we use absolute differences, in line with our fractional response approach, for which the *level* of the poverty rate is the natural metric.

The first brackets on the right-hand side contain the 'growth component', the second the 'distributional component' of poverty reduction. We refer to them as Y and D, respectively. Two of the four terms inside the brackets are counterfactuals: $H(\bar{y}_t/z, L_{t-1})$ is the poverty rate if the Lorenz curve had remained unchanged while average income evolved; $H(\bar{y}_{t-1}/z, L_t)$ is the poverty rate if average income had been kept constant while the distribution shifted. The decomposition is not exact. It is 'path dependent' and generates a residual.[28] In the absence of micro-data it is natural to approximate $H(\bar{y}_t/z, L_t)$ by our earlier function $H(\bar{y}_t/z, G_t)$, and replace the unknown quantities by predicted counterparts $\hat{H}(\bar{y}_t/z, G_{t-1})$ and $\hat{H}(\bar{y}_{t-1}/z, G_t)$.

We carry out this decomposition in three different data sets: one using the longest possible spells between surveys of the same type in 1981-2010; and two for the sub-periods 1981-2000 and 2000-2010. Since the availability of surveys before 2000 is more limited, the average spell length in the two sub-periods is comparable; both have a mean and median duration between the initial and final surveys of about seven years. The break point is not chosen by accident, but was picked by a formal structural change test. The turn of the millennium marked a qualitative change in the growth performance of Sub-Saharan Africa. For the next decade, the developing world collectively grew faster than the developed world.[29] The question is how much the poor benefited.

[27]Figure A-1 in Online Appendix A illustrates the latter.

[28]To get rid of the path dependency eq. (16) may be averaged with a second decomposition swapping t and $t-1$, delivering what Shorrocks (2013) calls the 'Shapley decomposition'. This is not attractive here because the decomposition residual also includes prediction errors.

[29]Online Appendix A discusses this in more detail, see Figure A-4 and Table A-3 in particular.

[Table 5 about here]

Following Kraay (2006), we borrow from the growth accounting literature to define the shares of growth and inequality (see e.g. Klenow and Rodriguez-Clare, 1997, Caselli, 2005). Within each of our three sample periods we compute the variances of the growth and distribution components, respectively $\mathrm{Var}(Y)$ and $\mathrm{Var}(D)$, and their covariance, $\mathrm{Cov}(Y, D)$. From these we obtain the share of the growth component as $s_Y = [\mathrm{Var}(Y) + \mathrm{Cov}(Y, D)]/[\mathrm{Var}(Y) + \mathrm{Var}(D) + 2\mathrm{Cov}(Y, D)]$, and the share of the distribution component as $s_D = 1 - s_Y$.[30] Table 5 reports the variance decomposition of all spells occurring within a particular region and period using the $2 a day poverty line. We also report the root mean squared error ($\sqrt{MSE}$) of the observed versus the predicted values and the number of spells for each region and period.

For all developing countries over the entire 1981-2010 period we find that the share of growth is about 82%, in line with the stylized facts from the literature. The covariance between the growth and distribution components is virtually zero, consistent with distribution-neutral growth. The residual is less than three quarters of a percentage point, indicating that even with added model error the decomposition works well.

Breaking down these results by region delivers more interesting insights. Growth dominated poverty reduction in Sub-Saharan Africa (where average incomes first declined, stagnated and then recovered again), in Latin America (where inequality was persistently high for most of the period in question), and in Europe and Central Asia (where the post-communist transition was accompanied by rapid movements in average incomes). Distributional change played a larger role in East Asia and Pacific, Middle East and North Africa, and South Asia, where it accounted for about one, two and three quarters of the poverty evolution, respectively. Some of these results come with a caveat, though. In South Asia, even though all countries but one are represented, the sample size is very small.[31] In the Middle East and North Africa, where survey coverage is notoriously bad, the residual is more than twice the average in size.

A new fact emerges once we split the sample at the turn of the millennium. Table 5 shows that, while growth was responsible for nearly all of the changes in the poverty rate before 2000, its share fell to 72% in the following decade. Since the mean spell lengths are about the same, this is not due to a long run vs. short run difference. Instead, this evolution seems to be the result of a shift in the poverty reducing quality of growth. Panel (b) shows that before 2000 the correlation between the growth and distribution component was negative for all developing countries as a whole and for half of the regions. Panel (c) demonstrates that after 2000 the correlation between the two

[30] One share can exceed 1, but the two shares always add up to 1. Caselli (2005) discusses a set of alternative measures, but they do not change the qualitative implications of our results.

[31] The large share of the distribution component is driven by the Maldives and Bhutan, where nearly all or about half of poverty reduction can be attributed to changes in distribution.

components turned positive for the sample as a whole and in all regions except East Asia. Online Appendix C confirms that this pattern also holds at the $1.25 a day poverty line.

[Figure 3 about here]

Figure 3 helps to make this point more tangible. The upper panel shows how the correlation between the estimated growth and distribution components changed over the turn of the millennium. Before 2000, reductions in poverty through growth often coincided with adverse distributional change (lower right quadrant), and negative economic growth often coincided with pro-poor distributional change (upper left quadrant), seemingly reflecting an Okunian trade-off. Since 2000, this is no longer true. In a majority of cases, poverty reduction through growth was reinforced by pro-poor distributional change (lower left quadrant). In other cases, growth was not accompanied by pro-poor redistribution (lower right quadrant). Only in a handful of cases did growth fail to reduce poverty (both upper quadrants). The lower panel of the figure shows that the estimated growth components closely track the observed change in the poverty rate in both subsamples, whether the change is positive or negative. We observe that, after 2000, growth becomes a net contributor to poverty *reduction* almost everywhere.

Taken together, these results challenge both Okun's 'big trade-off' and the view that growth is distribution-neutral. The evidence suggests that in developing countries growth can be, and has in fact become, more pro-poor since the turn of the millennium. It has done so both in the relative sense of benefiting the poor disproportionately, and in the absolute sense of reducing poverty rates. Whether this trend can be maintained is of formidable importance for the goal of ending absolute poverty over the coming decades.

5.4 Projections

In 2015, all member states of the United Nations agreed to 'eradicate extreme poverty for all people everywhere, currently measured as people living on less than $1.25 a day' by 2030. The key question we ask in this subsection is whether this new goal is attainable, or whether the global goalpost has been set so high that it may fail to promote a sensible allocation of development assistance.

The official World Bank regional poverty figures involve a considerable amount of interpolation and extrapolation, as most household surveys are carried out infrequently (for details see Chen and Ravallion, 2004). Our method is similar in spirit. It involves four steps. First, we extrapolate the last available survey income to 2010 using actually observed country growth rates in personal consumption expenditures per capita (PCE_{it}) from the national accounts.[32] Second, we project mean income into the future, using one

[32] The term 'national accounts' refers to data from the World Development Indicators or the Penn World Table 7.1, whichever has more data over the 30 year horizon.

of three constant-growth scenarios and one of three distributional scenarios. Third, we predict each country's poverty rate over the period from 2010 to 2030 using the estimates from column (2) in Table 2.[33] We may neglect measurement errors in income or inequality for forecasting purposes, implicitly assuming that their influence remains stable over time. Fourth, we calculate regional and global aggregates as population-weighted averages of the country estimates using World Bank population projections.

We base our constant-growth scenarios on the historical PCE_{it} growth rates in three predefined periods.[34] An 'optimistic' scenario uses the average PCE_{it} growth rate of each country over 2000-2010, a decade characterized by fast growth. A 'moderate' growth scenario uses the average PCE_{it} growth rate of each country over 1980-2010, the long-run average over the entire sample period. Finally, a 'pessimistic' growth scenario uses the 1980-2000 growth rates, a period during which mean consumption in Sub-Saharan Africa shrank at a rate of about 0.82% per year.

For each growth scenario, we simulate three inequality patterns. The 'distribution-neutral' scenario keeps inequality constant at its 2010 level, in line with the zero correlation between inequality changes and income growth observed over the entire period. In the 'pro-poor' scenarios the Gini coefficient declines, and in the 'pro-rich' scenarios it increases, by half a percent annually. Consider the following example for illustration. If a country's Gini coefficient is 0.40 in 2010 and growth is pro-poor, then by 2030 we project a Gini coefficient of about 0.36; but if growth is pro-rich, then the Gini coefficient rises to about 0.44 in 2030. Changes of this magnitude are in line with the population-weighted regional trends obtained from the surveys.

Our in-sample estimates for 2010 compare well with the 'official' World Bank figures. We can almost perfectly match the World Bank's results for the developing world total. We estimate a poverty rate of 40.37% in 2010 at the $2 a day poverty line, whereas the World Bank reports 40.67%.[35] Using the same population data, our estimates imply about 2.378 billion people under the $2 line versus 2.395 billion as reported by the World Bank. For three regions, our estimates are within one percentage point of the official figures; for the other three, they are within 3 percentage points.[36] In Online Appendix B we also discuss the results of two cross-validation exercises.[37] Our model very accurately forecasts regional or developing world aggregates and still has a decent out-of-sample

[33] Although the estimation only uses the sub-sample with $T_i \geq 2$, we can use the estimates to predict poverty for the entire sample ($T_i \geq 1$). We only lack estimates of the panel size effects for $T_i = 1$, so we assign these observations to the adjacent group ($T_i = 2$).

[34] Table A-3 in Online Appendix A reports population-weighted regional growth rates over these periods to illustrate the implied regional income dynamics.

[35] This was the official number until the Oct. 9, 2014 update of *PovcalNet*.

[36] For Sub-Saharan Africa, for instance, we estimate a poverty rate of 69.36% and the World Bank reports 69.87%. For East Asia, we estimate 26.78%, whereas the World Bank figure is 29.73%.

[37] We are indebted to an anonymous referee for suggesting this to us. Table B-4 in Online Appendix B uses the data before 2005 to predict regional poverty rates in the 2005–2010 period. Table B-5 conducts a 10-fold cross-validation where 10% of the time periods are deleted in each fold.

performance when it comes to individual observations.

[Table 6 about here]

Table 6 shows our projections for 2030 under the $2 poverty line. Our moderate growth scenario predicts that in 2030 about 1.87 billion people (26%) will live on less than $2 a day, as compared to 2.4 billion (40.37%) in 2010. However, much greater gains are possible. Global poverty under the $2 line falls by half to less than 20% of the developing world's population in the optimistic growth scenario with distribution-neutral or pro-poor growth. If this happens by 2030, then more than one billion people will have left poverty at the $2 line – undeniably a remarkable achievement.

Examining the regional distribution, we find that $2 poverty in East Asia is likely to fall to around 5% by 2030, down from 29.7% in 2010. Nearly everyone in East Asia will have entered the middle class by developing-country standards. Progress in South Asia is also likely to be rapid. According to our moderate growth estimate the expected poverty rate is 35.9% in 2030, meaning about 716 million poor, down from 66.7% or 1.1 billion poor in 2010. In the optimistic pro-poor growth case, the headcount ratio falls further to less than 20% and the number of poor to less than 400 million. In stark contrast, the $2 a day poverty rate in Sub-Saharan Africa is expected to remain very high. Our moderate growth scenario predicts a poverty rate of about 66%, down from 69.9% in 2010, which at current population projections implies almost one billion poor in Sub-Saharan Africa alone. Even in the optimistic and pro-poor growth scenario, we project a poverty rate of over 50% and more than 700 million poor. The bulk of the consumption distribution is too far below the $2 a day poverty line in 2010 for most of the subcontinent. Hence, poverty alleviation in Sub-Saharan Africa remains the primary development challenge of the first half of the 21st century.

These observations can in part be explained by a process of 'bunching up above $1.25 a day and just below $2 a day' occurring in East Asia and, to a lesser extent, in South Asia over the last two decades (Chen and Ravallion, 2010). These two regions have a relatively large population near the poverty line and hence most of the advances are projected to occur there. Latin America and the Caribbean, as well as the Middle East and North Africa, are richer and require stronger income growth to continuously reduce poverty. Sub-Saharan Africa, on the other hand, has a large proportion of poor far below the $2 a day line in 2010 (and almost 50% below $1.25 a day). It is facing a lower income elasticity and an income semi-elasticity that is below its peak. Hence, the subcontinent needs exceptionally strong income growth to make significant strides against poverty.

In Online Appendix B we repeat this analysis for the $1.25 a day poverty line. The key difference is that the pace of poverty reduction will have slowed considerably by 2030, both in terms of relative changes and in terms of numbers of poor people. Even under optimistic, distribution-neutral growth, we find a global poverty rate of 9.11% in 2030, with about 655 million people remaining extremely poor. Around 70% of the world's poor will live in Sub-Saharan Africa. There is a straightforward implication for the post-2015 development agenda: 2030 is unlikely to mark the end of extreme poverty, even under very optimistic assumptions. Figure 4 illustrates this result. It plots the historical evolution of world poverty from 1981 to 2010, a trend fitted through the observed data and extrapolated until 2030, and our different scenarios. The linear trend serves as a reference for the non-linear projections.

[Figure 4 about here]

Several points are worth noting. First, only the linear extrapolation which disregards the changing semi-elasticities predicts an extreme poverty rate in the vicinity of zero by 2030. Second, at the $1.25 a day line, all projections show a decelerating trend in poverty reduction, although in the optimistic scenarios the slowdown becomes noticeable later. Third, most scenarios suggest a $1.25 poverty rate higher than 10% in 2030; the optimistic pro-poor and distribution-neutral scenarios project 7.9% and 9.1%, respectively. Fourth, most scenarios also exhibit a decelerating pace of poverty reduction at the $2 a day line, but the slowdown tends to occur relatively late.

The changing composition of global poverty has profound implications for the medium-term future. Going forward, fast growing East Asia will contribute less and less to global poverty reduction, especially at the $1.25 a day line, while the share of the global poor residing in Sub-Saharan Africa and South Asia will continue to rise. None of our nine scenarios predict an extreme poverty rate near zero by 2030. This stands in stark contrast to earlier studies and the '3% by 2030' target of the World Bank which became enshrined in the Sustainable Development Goals of the UN. Such an outcome would require an implausible acceleration of growth in many of the poorest countries in the world, coupled with a deceleration of population growth. Even if growth rates in Sub-Saharan Africa were to *double* relative to the post-2000 trend, given persistent population growth, the global extreme poverty rate would still be well above 3% in 2030.

Two caveats are in order when it comes to comparing our estimates to recent figures from the World Bank. First, the World Bank switched from reporting global poverty in percent of the developing world population to percent of the global population. This lowered their estimates by a few percentage points and made the goal seem more attainable. Second, our results should not be very sensitive to the new 2011 PPPs, provided we shift the international poverty lines accordingly to $1.90 and $3.10 a day. The new poverty lines were re-drawn with the explicit aim of keeping the global yardstick constant (see Ferreira et al., 2015). The 2011 PPPs suggest that the poorest regions are a little less poor and middle-income regions a little poorer, but the global poverty rate in 2010 only changed by about one tenth of a percentage point. This is trivial in comparison with the fundamental uncertainties involved in estimating such a global number.

6 Concluding remarks

The main purpose of this paper is to provide an assessment of the potential impacts and historical contributions of growth and redistribution to poverty reduction which reflects the heterogeneity of initial conditions in the developing world. To this end, we develop a new empirical framework for poverty accounting. Our approach has several advantages over the linear models with interactions typically employed in the literature. Above all, it accommodates the fact that the poverty headcount ratio is a fraction, and it closely approximates the shape of the Lorenz curve near the poverty line using only two summary statistics: mean income and the Gini coefficient. It is also general enough to be applicable in other contexts than global poverty. It could, for instance, be used to decompose national poverty rates across geographical areas or population groups.

We use this new framework together with a large data set of aggregated household surveys to show that differences in initial income and inequality across countries induce strong regional heterogeneity in the estimated poverty responses. Unlike linear approaches, our model precisely estimates the non-linear semi-elasticities and elasticities of poverty across the entire range of the data, thus allowing policy makers to understand their individual circumstances. We emphasize that the semi-elasticities are the policy relevant parameters, assuming that policy makers care about differences in the share of the population that is poor, rather than relative changes in the poor population.

We find that the relative importance of growth and inequality has changed since the turn of the millennium. Growth has become more pro-poor, not only in the sense that it actually reduces poverty, but also in the sense that it is often accompanied by lower inequality. As a result, inequality has begun to play a greater role. The share of poverty reduction due to changes in distribution increased from less than 10% before 2000 to about a third thereafter. To better understand the implications of this development, we project global and regional poverty rates from 2010 until 2030 using a variety of scenarios. A billion people may be lifted out of $2 a day poverty by 2030, but progress at the $1.25 a day line is bound to slow down. Simultaneously, the global geography of poverty is shifting. There are two major consequences. First, the goal of ending extreme poverty within a generation will not be achieved. Second, global poverty will be concentrated in Sub-Saharan Africa, especially in states that are currently fragile and conflict-ridden.

Finally, while pro-poor redistribution diminishes poverty by definition, it can also make economic growth much more effective at reducing poverty. We argue that this interaction between inequality and growth is of great importance and reverses some long-held views, as far as the developing world is concerned. At the same time, it should be remembered that the scope of poverty accounting is limited to the proximate determinants of poverty. It does not identify the ultimate causes of poverty, or how growth and inequality are causally related, but focuses on the last link in a long chain of events.

References

Aaron, H. (1967). The foundations of the "war on poverty" reexamined. *American Economic Review 57*(5), 1229–1240.

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

Aitchison, J. and J. A. C. Brown (1957). *The Lognormal Distribution: With Special References to Its Uses in Economics*. Cambridge University Press.

Besley, T. and R. Burgess (2003). Halving global poverty. *Journal of Economic Perspectives 17*(3), 3–22.

Blundell, R. W. and J. L. Powell (2004). Endogeneity in semiparametric binary response models. *Review of Economic Studies 71*(3), 655–679.

Bourguignon, F. (2003). The growth elasticity of poverty reduction: Explaining heterogeneity across countries and time periods. In T. S. Eicher and S. J. Turnovsky (Eds.), *Inequality and Growth: Theory and Policy Implications*, pp. 3–26. Cambridge, MA: MIT Press.

Bresson, F. (2009). On the estimation of growth and inequality elasticities of poverty with grouped data. *Review of Income and Wealth 55*(2), 266–302.

Caselli, F. (2005). Accounting for cross-country income differences. Volume 1, Part A of *Handbook of Economic Growth*, Chapter 9, pp. 679–741. Elsevier.

Chamberlain, G. (1984). Panel data. Volume 2 of *Handbook of Econometrics*, pp. 1247–1318. Elsevier.

Chambers, D. and S. Dhongde (2011). A non-parametric measure of poverty elasticity. *Review of Income and Wealth 57*(4), 683–703.

Chen, S. and M. Ravallion (2004). How have the world's poorest fared since the early 1980s? *The World Bank Research Observer 19*(2), 141–169.

Chen, S. and M. Ravallion (2010). The developing world is poorer than we thought, but no less successful in the fight against poverty. *Quarterly Journal of Economics 125*(4), 1577–1625.

Chesher, A. and C. Schluter (2002). Welfare measurement and measurement error. *Review of Economic Studies 69*(2), 357–378.

Chotikapanich, D., W. E. Griffiths, and D. S. P. Rao (2007). Estimating and combining national income distributions using limited data. *Journal of Business & Economic Statistics 25*(1), 97–109.

Datt, G. and M. Ravallion (1992). Growth and redistribution components of changes in poverty measures: A decomposition with applications to Brazil and India in the 1980s. *Journal of Development Economics 38*(2), 275–295.

Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics 87*(1), 1–19.

Dollar, D., T. Kleineberg, and A. Kraay (2016). Growth still is good for the poor. *European Economic Review 81*, 68–85.

Dollar, D. and A. Kraay (2002). Growth is good for the poor. *Journal of Economic Growth 7*(3), 195–225.

Ferreira, F. H., S. Chen, Y. Dikhanov, N. Hamadeh, D. Jolliffe, A. Narayan, E. B. Prydz, A. Revenga, P. Sangraula, U. Serajuddin, and N. Yoshida (2015, October). A global count of the extreme poor in 2012: data issues, methodology and initial results. Policy Research Working Paper 7432, The World Bank.

Gourieroux, C., A. Monfort, and A. Trognon (1984). Pseudo maximum likelihood methods: Theory. *Econometrica 52*(3), 681–700.

Hausman, J. A. (1978). Specification tests in econometrics. *Econometrica 46*(6), 1251–1271.

Heckman, J. J. and R. Robb (1985). Alternative methods for evaluating the impact of interventions: An overview. *Journal of Econometrics 30*(1), 239–267.

Hoover, G. A., W. Enders, and D. G. Freeman (2008). Non-white poverty and macroeconomy: The impact of growth. *American Economic Review 98*(2), 398–402.

Kakwani, N. (1993). Poverty and economic growth with application to Côte d'Ivoire. *Review of Income and Wealth 39*(2), 121–139.

Kalwij, A. and A. Verschoor (2007). Not by growth alone: The role of the distribution of income in regional diversity in poverty reduction. *European Economic Review 51*(4), 805–829.

Klasen, S. and M. Misselhorn (2008). Determinants of the growth semi-elasticity of poverty reduction. Working paper, Ibero America Institute for Economic Research.

Klenow, P. and A. Rodriguez-Clare (1997). The neoclassical revival in growth economics: Has it gone too far? In B. S. Bernanke and J. J. Rotemberg (Eds.), *NBER Macroeconomics Annual 1997*, Volume 12, pp. 73–103. Cambridge, MA: MIT Press.

Kraay, A. (2006). When is growth pro-poor? Evidence from a panel of countries. *Journal of Development Economics 80*(1), 198–227.

Kraay, A., T. Kleineberg, and D. Dollar (2014). Growth, inequality, and social welfare: Cross-country evidence. Working paper, Centre for Economic Policy Research.

Krause, M. (2014). Parametric lorenz curves and the modality of the income density function. *Review of Income and Wealth 60*(4), 905–929.

Lin, W. and J. M. Wooldridge (2017, March). Testing and correcting for endogeneity in nonlinear unobserved effects models. Manuscript.

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica 46*(1), 69–85.

Papke, L. E. and J. M. Wooldridge (1996). Econometric methods for fractional response variables with an application to 401(k) plan participation rates. *Journal of Applied Econometrics 11*(6), 619–632.

Papke, L. E. and J. M. Wooldridge (2008). Panel data methods for fractional response variables with an application to test pass rates. *Journal of Econometrics 145*(1–2), 121–133.

Ravallion, M. (2001). Growth, inequality and poverty: Looking beyond averages. *World Development 29*(11), 1803–1815.

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

Ravallion, M. and S. Chen (1997). What can new survey data tell us about recent changes in distribution and poverty? *World Bank Economic Review 11*(2), 357–382.

Ravallion, M. and S. Chen (2003). Measuring pro-poor growth. *Economics Letters 78*(1), 93–99.

Santos Silva, J. and S. Tenreyro (2006). The log of gravity. *Review of Economics and Statistics 88*(4), 641–658.

Shorrocks, A. F. (2013). Decomposition procedures for distributional analysis: a unified framework based on the Shapley value. *Journal of Economic Inequality 11*(1), 99–126.

Wooldridge, J. M. (2010a, May). Correlated random effects models with unbalanced panels. Manuscript.

Wooldridge, J. M. (2010b). *Econometric Analysis of Cross Section and Panel Data* (2nd ed.). The MIT Press.

World Bank (2005). World development report 2006: Equity and development. Technical report, The World Bank.

World Bank (2015). A measured approach to ending poverty and boosting shared prosperity. Technical report, The World Bank.

Young, A. (2011, December). The Gini coefficient for a mixture of ln-normal populations. Manuscript. The London School of Economics and Political Science, London, UK.

Tables and Figures

Table 1 – Summary statistics

	Mean	Std. Deviation	Min	Max	N
Main variables (2005 PPP $)					
H_{it} – Headcount ($2)	0.303	0.286	.0002	.9845	809
H_{it} – Headcount ($1.25)	0.182	0.219	.0002	.9255	789
G_{it} – Gini coefficient	0.424	0.102	.2096	.7433	809
$\bar{y}_{it}$ – Mean income or expenditure in $ per month	194.59	125.90	14.93	766.78	809
PCE_{it}^{P} – Consumption (PWT) in $ per month	338.64	234.59	14.39	1231.21	795
Survey type dummies					
Consumption expenditures, grouped	0.611	0.488	0	1	809
Consumption expenditures, household	0.015	0.121	0	1	809
Income, grouped	0.132	0.339	0	1	809
Income, household	0.242	0.429	0	1	809

Notes: The table reports summary statistics. All aggregate survey data originate from the World Bank's *PovcalNet* database. National accounts consumption has been calculated based on the Penn World Table 7.1.

Table 2 – Fractional probit models (QML) – Dependent variable: H_{it}, \$2 a day

	(1) Regular CRE		(2) Unbalanced		(3) Unbalanced + Two-Step	
	H_{it}	APEs	H_{it}	APEs	H_{it}	APEs
$\ln \bar{y}_{it}$	-1.263	-0.284	-0.880	-0.281	-1.086	-0.338
	(0.054)	(0.012)	(0.048)	(0.011)	(0.190)	(0.035)
$\ln G_{it}$	1.032	0.232	0.786	0.251	0.824	0.256
	(0.114)	(0.026)	(0.098)	(0.026)	(0.160)	(0.030)
$\hat{\nu}_{it}$					0.200	
					(0.112)	
CRE (Corr. Rand. Effects)	Yes		Yes		Yes	
Survey type dummies	Yes		Yes		Yes	
Time dummies	Yes		Yes		Yes	
Panel size dummies	No		Yes		Yes	
Panel size dummies × CRE	No		Yes		Yes	
Variance function	No		Yes		Yes	
Scale factor	0.225		0.319		0.311	
$N \times \bar{T}$	789		789		775	
N	104		104		103	
pseudo R^2	0.984		0.993		0.994	
$\ln \mathcal{L}$	-219.3		-315.6		-313.2	
$\sqrt{MSE}$	0.0355		0.0238		0.0223	

Notes: The table reports fractional response QML estimates. The dependent variable is the poverty rate at \$2 a day (in 2005 PPPs). 20 observations with $T_i = 1$ are not used in estimation. The panel structure is country-survey-year. In models (1) and (2), the standard errors of the coefficients are robust to clustering at the country level and the standard errors of the APEs are computed via the delta method. We include the time averages of the survey type and time dummies in (2) and (3), but constrain their coefficients to be equal across the panel sizes. The standard errors of the coefficients and the APEs in model (3) account for the first stage estimation step with a panel bootstrap using 999 bootstrap replications. The linear projection in the first stage uses $\ln PCE^P_{it}$ as an instrument for $\ln \bar{y}_{it}$. The first-stage cluster-robust F-statistic for instrument relevance in (3) is 28.05. The estimation of model (3) excludes West Bank and Gaza entirely (2 observations) and 12 observations from ECA countries pre-1990 for lack of PCE data.

Table 3 – Income elasticities and semi-elasticities, \$2 a day, by region

	Time period				
	1981–1989	1990–1994	1995–1999	2000–2004	2005–2010
	Panel (a) Regional income elasticities				
East Asia and Pacific	-0.975	-1.012	-1.217	-1.120	-1.556
	(0.031)	(0.033)	(0.054)	(0.042)	(0.101)
Eastern Europe and Central Asia	-4.346	-2.872	-2.681	-2.824	-3.282
	(0.575)	(0.318)	(0.284)	(0.312)	(0.396)
Latin America and Caribbean	-2.249	-2.339	-2.388	-2.312	-2.947
	(0.247)	(0.261)	(0.276)	(0.263)	(0.376)
Middle East and North Africa	-2.149	-2.091	-2.002	-1.943	-2.480
	(0.206)	(0.191)	(0.170)	(0.163)	(0.252)
South Asia	-0.541	-0.621	-0.796	-1.008	-1.177
	(0.055)	(0.050)	(0.032)	(0.032)	(0.046)
Sub-Saharan Africa	-0.815	-0.423	-0.424	-0.577	-0.616
	(0.029)	(0.041)	(0.042)	(0.036)	(0.035)
	Panel (b) Regional income semi-elasticities				
East Asia and Pacific	-0.565	-0.571	-0.586	-0.582	-0.556
	(0.033)	(0.036)	(0.046)	(0.042)	(0.051)
Eastern Europe and Central Asia	-0.032	-0.218	-0.265	-0.229	-0.137
	(0.008)	(0.015)	(0.020)	(0.015)	(0.010)
Latin America and Caribbean	-0.384	-0.358	-0.344	-0.366	-0.202
	(0.027)	(0.024)	(0.023)	(0.025)	(0.013)
Middle East and North Africa	-0.412	-0.429	-0.453	-0.469	-0.318
	(0.034)	(0.037)	(0.042)	(0.043)	(0.025)
South Asia	-0.415	-0.454	-0.522	-0.570	-0.585
	(0.022)	(0.018)	(0.022)	(0.035)	(0.044)
Sub-Saharan Africa	-0.528	-0.346	-0.346	-0.433	-0.452
	(0.023)	(0.020)	(0.020)	(0.015)	(0.015)

Notes: The table reports regional income elasticities in panel (a) and semi-elasticities in panel (b). The estimates are computed by plugging period and region-specific averages of mean income and inequality into eq. (15) or its semi-elasticity counterpart, and then averaging over the entire sample. Standard errors are obtained via a panel bootstrap using 999 replications.

Table 4 – Inequality elasticities and semi-elasticities, \$2 a day, by region

	Time period				
	1981–1989	1990–1994	1995–1999	2000–2004	2005–2010
	Panel (a) Regional inequality elasticities				
East Asia and Pacific	0.740	0.768	0.924	0.850	1.181
	(0.100)	(0.097)	(0.108)	(0.103)	(0.138)
Eastern Europe and Central Asia	3.298	2.180	2.034	2.143	2.491
	(0.505)	(0.301)	(0.277)	(0.291)	(0.348)
Latin America and Caribbean	1.707	1.775	1.812	1.755	2.236
	(0.182)	(0.194)	(0.196)	(0.186)	(0.266)
Middle East and North Africa	1.631	1.587	1.520	1.475	1.882
	(0.192)	(0.192)	(0.190)	(0.179)	(0.247)
South Asia	0.411	0.471	0.604	0.765	0.893
	(0.089)	(0.092)	(0.090)	(0.102)	(0.121)
Sub-Saharan Africa	0.618	0.321	0.322	0.438	0.468
	(0.083)	(0.054)	(0.058)	(0.064)	(0.066)
	Panel (b) Regional inequality semi-elasticities				
East Asia and Pacific	0.429	0.433	0.444	0.442	0.422
	(0.051)	(0.051)	(0.052)	(0.051)	(0.050)
Eastern Europe and Central Asia	0.024	0.166	0.201	0.174	0.104
	(0.007)	(0.015)	(0.018)	(0.017)	(0.013)
Latin America and Caribbean	0.291	0.272	0.261	0.277	0.153
	(0.044)	(0.041)	(0.042)	(0.044)	(0.029)
Middle East and North Africa	0.313	0.325	0.344	0.356	0.242
	(0.039)	(0.038)	(0.039)	(0.041)	(0.024)
South Asia	0.315	0.345	0.396	0.433	0.444
	(0.056)	(0.055)	(0.050)	(0.051)	(0.052)
Sub-Saharan Africa	0.400	0.262	0.263	0.328	0.343
	(0.048)	(0.038)	(0.041)	(0.042)	(0.042)

Notes: The table reports regional inequality elasticities in panel (a) and semi-elasticities in panel (b). The estimates are computed by plugging period and region-specific averages of mean income and inequality into eq. (15) or its semi-elasticity counterpart, and then averaging over the entire sample. Standard errors are obtained via a panel bootstrap using 999 replications.

Table 5 – Decomposition at $2 a day poverty line, by region

	Var(Y)	Var(D)	Cov(Y, D)	s_Y	s_D	$\sqrt{MSE}$	N
	Panel (a) Spells from 1981 to 2010						
East Asia and Pacific	1.166	0.385	0.001	75.16	24.84	1.01	12
Europe and Central Asia	3.429	0.414	0.164	86.14	13.86	0.57	39
Latin America and Caribbean	3.073	0.666	-0.619	98.12	1.88	0.79	28
Middle East and North Africa	0.425	0.452	0.304	49.10	50.90	1.38	8
South Asia	0.659	2.743	0.191	22.46	77.54	0.36	7
Sub-Saharan Africa	1.482	0.134	-0.196	105.08	-5.08	0.48	29
All developing	2.576	0.579	0.002	81.61	18.39	0.73	123
	Panel (b) Spells from 1981 to 1999						
East Asia and Pacific	0.951	0.151	0.017	85.23	14.77	0.86	9
Europe and Central Asia	14.568	1.912	-0.706	91.99	8.01	0.92	25
Latin America and Caribbean	3.750	1.229	-0.633	83.94	16.06	0.89	26
Middle East and North Africa	0.372	0.365	0.207	50.31	49.69	0.78	6
South Asia	0.244	0.021	0.004	91.11	8.89	0.33	4
Sub-Saharan Africa	2.217	0.702	-0.309	82.89	17.11	0.92	17
All developing	6.852	1.073	-0.368	90.19	9.81	0.87	87
	Panel (c) Spells from 2000 to 2010						
East Asia and Pacific	1.579	1.143	-0.134	58.89	41.11	1.18	10
Europe and Central Asia	3.544	0.498	0.826	76.75	23.25	0.57	26
Latin America and Caribbean	0.365	0.150	0.010	70.10	29.90	0.35	19
Middle East and North Africa	0.780	0.741	0.629	50.69	49.31	1.76	5
South Asia	0.642	1.122	0.731	42.57	57.43	0.40	6
Sub-Saharan Africa	2.051	0.560	0.188	74.97	25.03	0.87	21
All developing	2.049	0.559	0.365	72.31	27.69	0.81	87

Notes: The table reports the results of the decomposition of the observed changes in the poverty rate at $2 a day into its growth and distribution components at the regional level. The estimates are based on column (3) of Table 2. Panels (a) to (c) run this decomposition over the three different periods specified. We predict the counterfactual quantities using the first and last available observations from the longest run of surveys of the same type within the specified period.

Table 6 – Projected poverty headcount ratios and poor population at $2 a day in 2030, by region

	Average PCE Growth								
	Optimistic (2000-2010)			Moderate (1980-2010)			Pessimistic (1980-2000)		
	Change in Inequality (Gini)								
	pro-poor	neutral	pro-rich	pro-poor	neutral	pro-rich	pro-poor	neutral	pro-rich
Panel (a) – Headcount at $2 a day in 2030 (in percent)									
East Asia and Pacific	3.79	4.61	5.55	4.03	4.90	5.90	4.39	5.31	6.36
Europe and Central Asia	0.45	0.56	0.69	2.80	3.13	3.49	9.35	10.28	11.31
Latin America and Caribbean	4.00	4.73	5.59	6.29	7.39	8.66	8.57	9.99	11.60
Middle East and North Africa	2.85	3.55	4.39	7.20	8.62	10.25	12.86	14.88	17.12
South Asia	19.83	23.12	26.74	31.60	35.88	40.39	40.35	45.00	49.73
Sub-Saharan Africa	51.62	54.56	57.46	63.67	66.36	68.98	70.93	73.32	75.63
Total	17.36	19.23	21.24	23.73	25.94	28.26	28.71	31.07	33.53
Panel (b) – Poor population at $2 a day in 2030 (in millions)									
East Asia and Pacific	82.42	100.20	120.61	87.49	106.39	128.23	95.44	115.26	138.11
Europe and Central Asia	2.12	2.63	3.26	13.24	14.79	16.49	44.25	48.63	53.50
Latin America and Caribbean	28.39	33.63	39.73	44.71	52.52	61.50	60.90	70.99	82.42
Middle East and North Africa	12.64	15.74	19.47	31.94	38.22	45.46	57.02	65.95	75.91
South Asia	395.33	461.03	533.21	629.97	715.46	805.25	804.56	897.12	991.46
Sub-Saharan Africa	723.19	764.47	805.00	892.01	929.76	966.48	993.71	1027.19	1059.59
Total	1249.19	1383.61	1528.08	1707.20	1866.01	2033.35	2065.12	2235.43	2412.30

Notes: The table reports forecasts of the $2 a day poverty rate in 2030. The forecasts are based on the estimates reported in Column (2) of Table 2 and the different growth and distribution scenarios outlined in the text. Population projections are from the World Bank's Health, Nutrition and Population Statistics database. The survey data are from the World Bank's *PovcalNet* database.

Figure 1 – Transformed headcount ($2 a day) and log-mean income, by region

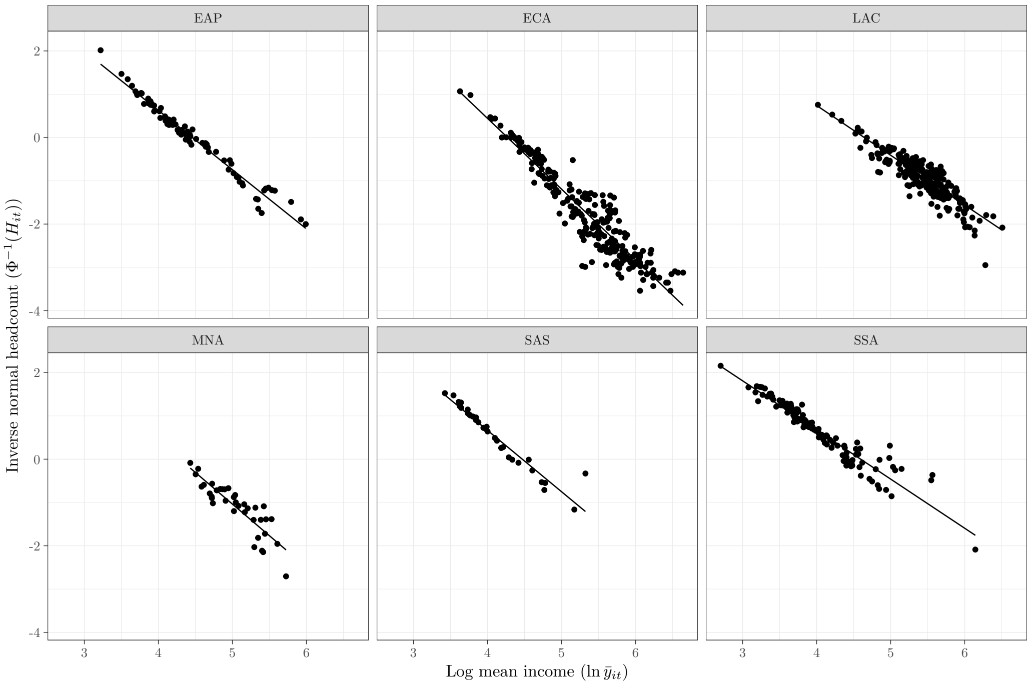

Notes: Illustration of how effectively the inverse normal cdf linearizes the conditional mean. The figure plots $\Phi^{-1}(H_{it})$ against $\ln \bar{y}_{it}$ per region, including a regression line.

Figure 2 – Predicted income and Gini elasticities and semi-elasticities of poverty, \$2 a day

(a) elasticities ($\hat{\varepsilon}^{H\bar{y}}$ and $\hat{\varepsilon}^{HG}$)

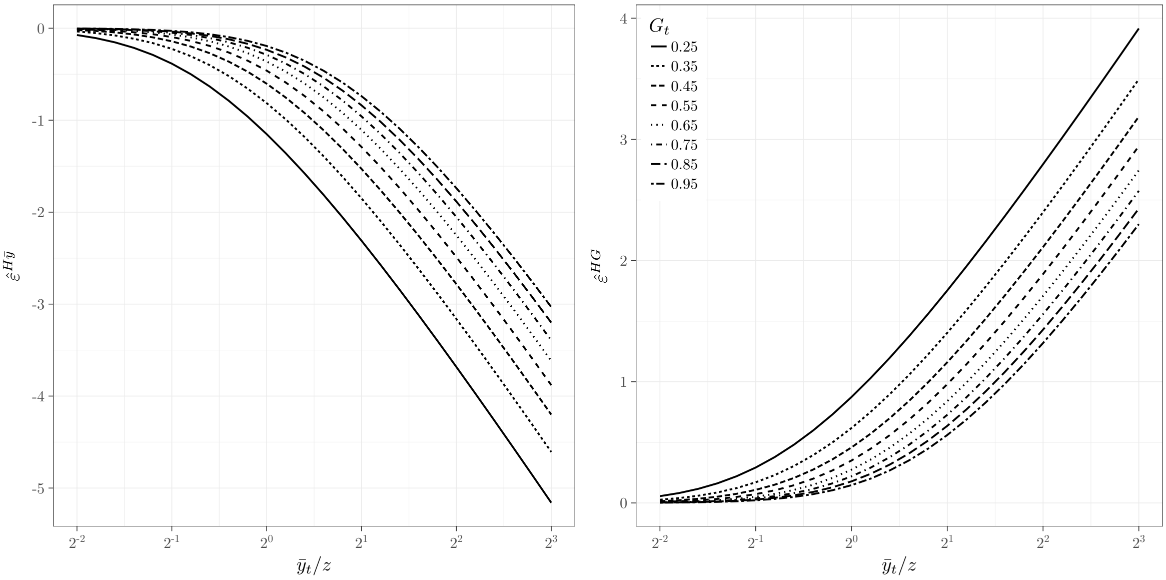

(b) semi-elasticities ($\hat{\eta}^{H\bar{y}}$ and $\hat{\eta}^{HG}$)

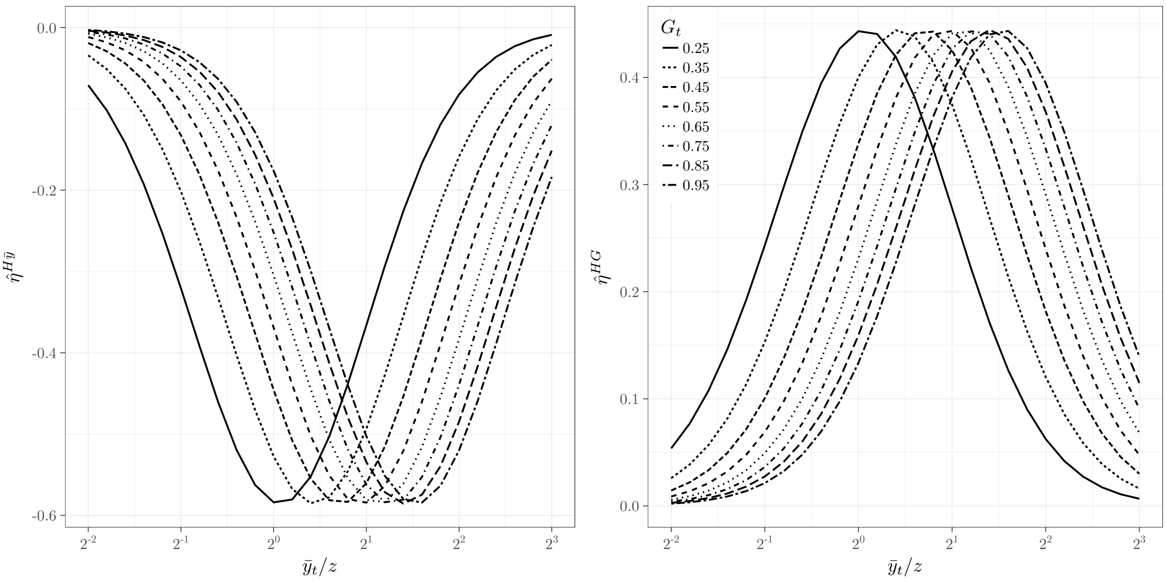

Notes: Illustration of the non-linear nature of the poverty-growth-inequality relationship based on the estimates presented in column (3) of Table 2. Panel (a) shows the estimated income elasticities of poverty (on the left) and estimated inequality elasticities of poverty (on the right) plotted over the ratio of average income to the poverty line. Panel (b) shows the estimated income semi-elasticities of poverty (on the left) and estimated inequality semi-elasticities of poverty (on the right) plotted over the ratio of average income to the poverty line. The various curves correspond to estimates based on different Gini coefficients.

Figure 3 – Estimated components from poverty accounting, \$2 a day

(a) Growth versus distribution component, pre and post 2000

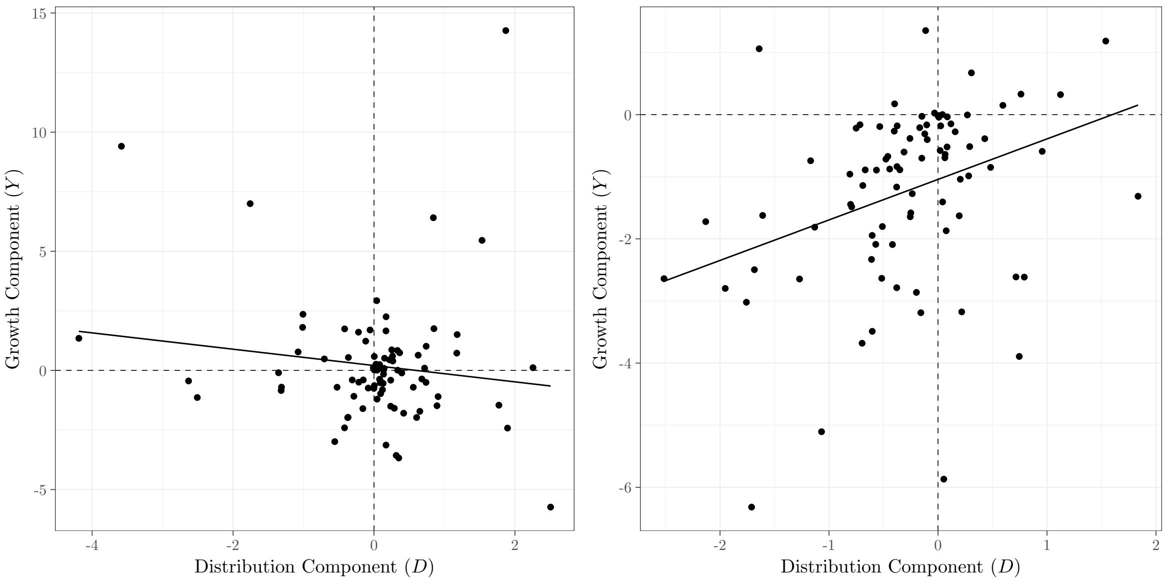

(b) Observed change versus growth component, pre and post 2000

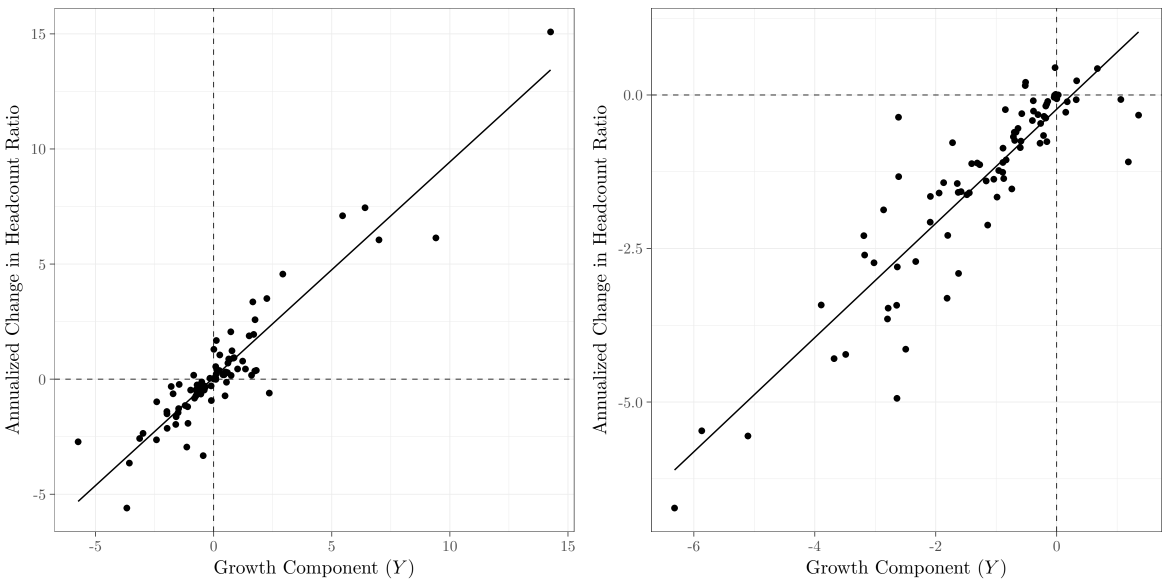

Notes: Illustration of the results from the poverty decomposition. Panel (a) plots the growth component against the distribution component for the pre-2000 sample (on the left) and for the post-2000 sample (on the right). Panel (b) plots the change in the poverty rate against the distribution component for the pre-2000 sample (on the left) and for the post-2000 sample (on the right). The number of dots in a graph may exceed the number of countries, since some countries have both a consumption and an income survey spell.

Figure 4 – Global poverty: observed data, trends and projections, \$1.25 and \$2 a day

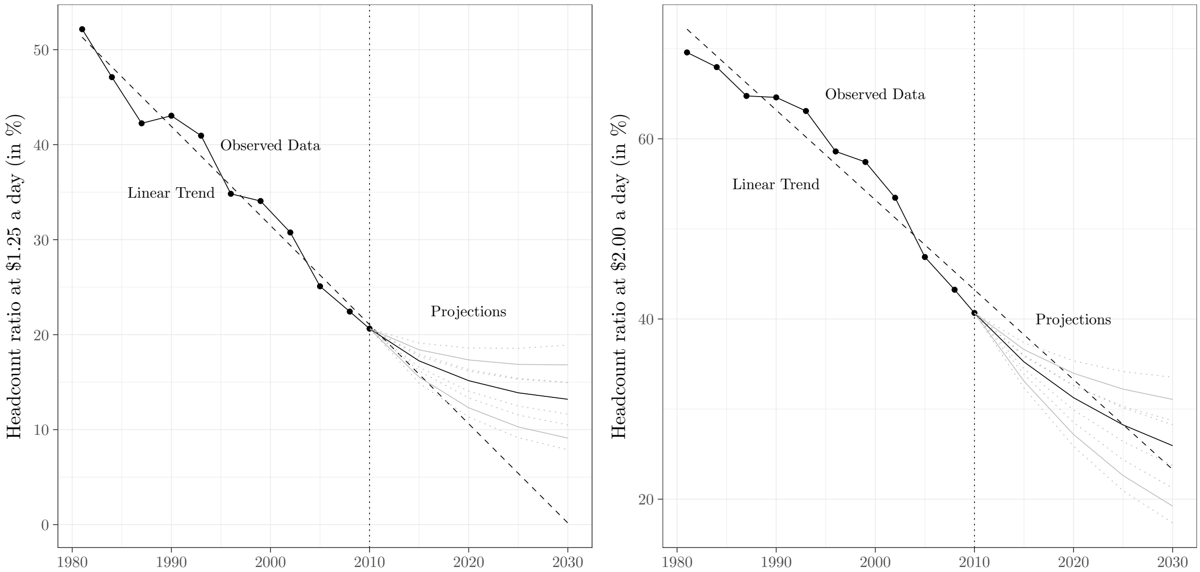

Notes: Illustration of the projections. The left panel shows the projections at the \$1.25 a day poverty line; the right panel shows the projections at \$2 a day poverty line. The dashed line is a linear trend based on the data until 2010. The black line after 2010 displays the moderate growth scenario, the two gray lines above and below stand for the pessimistic and optimistic growth scenarios with distribution-neutral growth, respectively. Small dashed curves above and below each growth scenario are the associated pro-rich and pro-poor scenarios.

Online appendix

A Additional summary statistics

A.1 Household surveys in the sample

Our sample is severely unbalanced and mixes different types of survey aggregates. Poverty and inequality were estimated by *PovcalNet* either on data that have been pre-aggregated into groups by the national statistical offices, or directly on the basis of the underlying household surveys.

Table A-1 reports the survey type by geographic region and illustrates the temporal coverage. ‘I’ and ‘C’ denote grouped income or consumption data, whereas ‘i’ and ‘c’ denote unit level income or consumption data. Note that Latin America is the only region which predominantly relies on income surveys. Table A-2 shows additional summary statistics by region.

The other figures complement statements made in the main text. Figure A-1 illustrates the lack of a raw correlation between changes in distribution and income growth in the sample as a whole. Figure A-2 shows that the inverse normal cdf also linearizes the relationship between the observed poverty rate and the log of the Gini coefficient. Note that this is the simple bivariate relationship without partialling out the influence of other variables, such as log income. Figure A-3 adds the regional trends in the raw poverty rates over time.

Table A-1 – Survey coverage and types, 1981 – 2010

	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10
Panel (a) East Asia and Pacific																														
CHN	I			I			I			C		C	C	C	C	C	C	C	C			C			C			C	C	
FJI																							C						C	
IDN				C			C			C			C			C			C			C			C	C	C	C	C	C
KHM														C										C			C	C	C	
LAO												C					C					C						C		
MYS				I			I		I			I			I		I							I			I		I	
PHL					C			C			C			C			C			C			C			C			C	
PNG																C														
THA	C							C		c		c		c		c		c	c	c		c				c		c		c
TMP																					C						C			
VNM													C					C				C		C		C		C		
Panel (b) Eastern Europe and Central Asia																														
ALB																	C					C		C	C			C		
ARM																I			C		C	C	C	C	C	C	C	C		C
AZE															C						C							C		
BGR									C					C	I		C				C		C				C			
BIH																					C			C			C			
BLR								I					I		I			I		C	C	C		C	C	C	C	C	C	C

Continued on next page

Table A-1 – *Continued from previous page*

	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10
CZE													I			I														
EST								I					I		C			C		C	C	C	C	C						
GEO																C	C	C	C	C	C	C	C		C	C	C	C	C	C
HRV								I										C	C	C	C			C				C		
HUN							I		I				I					C	C	C	C	C		C			C			
KAZ								I					I			C					C	C	C	C		C	C	C	C	
KGZ													C					C				C		C	C	C	C	C	C	C
LTU								I					I			C		C		C	C	C		C				C		
LVA								I					I		I	I	I	C					C	C			C	C	C	
MDA								I				I					C	C	C		C	C	C	C	C	C	C	C	C	C
MKD																		C		C		C	C	C	C	C		C	C	C
MNE																									C	C	C	C		
POL					I		I		I			C	I			C		C	C	C	C	C		C	C	C	C	C	C	C
ROM									I			I		I				C		C	C	C	C	C	C	C	C	C		C
RUS								I					C			C			C		C	C	C	C	C	C	C	C	C	
SRB																						C	C	C	C	C	C	C	C	C
SVK																I								C	C	C	C	C	C	
SVN							I						I					C				C	C	C						
TJK																			C				C	C			C		C	
TKM								I					I					C												
TUR							C							C								C	C	C	C	C	C	C	C	C
UKR								I							C	C			I			C	C	C	C	C	C	C	C	C
	Panel (c) Latin America and the Caribbean																													
ARG							I				I	I	I	I	I	I	I	I	I	I	I	I	I	I	I	I	I	I	I	I
BLZ													i	i	C	i	i	i	i											
BOL											I		I				I		I	I	I	I			I	I	I	I		
BRA	i	i	i	i	i	i	i	i	i	i		i	i		i	i	i	i	i		i	i	i	i	i	i	i	i	i	
CHL							i			i		i		i		i		i		i			i			i			i	
COL												i				i			i	i	i	i	i	i	i	i	i	i	i	i
CRI	I					I			i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	i	
DOM						I			I			I				i	i			i	i	i	i	i	i	i	i	i	i	i
ECU							I							i				i	i	i			i		i	i	i	i	i	i
GTM							I		I									I		i		i	i	i		i				
GUY													I					I												
HND									I	I	i	i	i	i	i	i	i	i	i		i	i	i	i	i	i	i	i	i	
HTI																					i									
JAM								C		C			C			C			C			C		C						
LCA															I															
MEX				C								C		C		C		C		C		C		C		C		C		C
NIC													C					C			C				C					
PAN									i		i				i		i	c			i	i	i	i	i	i			i	i
PER						C								C			i	i	i	i	i	i	i	i	i	i	i	i	i	i
PRY										i					i			i	i		i	i	i	i	i	i	i	i	i	i

Continued on next page

Table A-1 – *Continued from previous page*

	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10
SLV											i				i	i		i	i		i	i	i	i	i	i	i	i	i	
SUR																			I											
TTO								I				I																		
URY	I								i																	i	i	i	i	i
VEN	I						I		i			i			i			i	i		i	i	i	i	i	i				
Panel (d) Middle East and North Africa																														
DJI																						C								
DZA								C							C															
EGY											C					C				C					C			C		
IRN						C				C				C				C							C					
IRQ																											C			
JOR							C					C					C						C			C		C		C
MAR					C						C								C		C						C			
SYR																								C						
TUN					C					C					C					C					C					C
WBG																											C		C	
YEM																		C							C					
Panel (e) South Asia																														
BGD				C		C			C			C				C				C					C					C
BTN																							C				C			
IND			C					C						C											C					C
LKA					C						C					C						C					C			C
MDV																		C						C						
NPL					I											C							C							C
PAK							C				C						C		C			C			C	C		C		
Panel (f) Sub-Saharan Africa																														
AGO																				C									C	
BDI												C						C								C				
BEN																							C							
BFA														C				C					C						C	
BWA						C								C																
CAF												C											C					C		
CIV					C	C	C	C					C		C			C				C						C		
CMR																C					C						C			
COG																									C					
COM																								C						
CPV																						C								
ETH		C													C					C					C					
GAB																									C					
GHA								C	C			C						C								C				
GIN											C			C									C				C			
GMB																		C					C							
GNB													C									C								

Continued on next page

Table A-1 – *Continued from previous page*

	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10
KEN												C		C			C								C					
LBR																											C			
LSO							C						C	C									C							
MDG													C				C		C		C				C					C
MLI														C							C					C				C
MOZ																C							C					C		
MRT							C						C			C				C				C				C		
MWI																		C						C						C
NAM													I											I						
NER												C		C											C			C		
NGA						C						C				C								C						C
RWA					C															C						C				
SDN																													C	
SEN											C			C							C				C					
SLE																							C							
STP																					C									
SWZ															C						C									C
SYC																											C			
TCD																							C							
TGO																										C				
TZA												C								C							C			
UGA									C			C				C			C			C				C			C	
ZAF													C		C					C						C			C	
ZAR																										C				
ZMB													C			C		C					C	C		C				C

Notes: The table reports the temporal and regional coverage of the survey data, as well as the type of survey conducted in each year. ‘I’ and ‘C’ denote grouped income or consumption data, whereas ‘i’ and ‘c’ denote unit level income or consumption data.

Table A-2 – Summary statistics by region (unweighted)

Variable	Mean	Standard Deviation	Min	Max
	East Asia and Pacific (N=80)			
H_{it} – Headcount ($2)	0.502	0.267	0.023	0.978
G_{it} – Gini coefficient	0.392	0.058	0.275	0.509
$\bar{y}_{it}$ – Mean income or expenditure	107.86	78.39	25.02	399.76
	Eastern Europe and Central Asia (N=254)			
H_{it} – Headcount ($2)	0.110	0.169	0.000	0.857
G_{it} – Gini coefficient	0.330	0.056	0.210	0.537
$\bar{y}_{it}$ – Mean income or expenditure	251.99	136.11	37.66	766.78
	Latin America and Caribbean (N=274)			
H_{it} – Headcount ($2)	0.204	0.122	0.002	0.775
G_{it} – Gini coefficient	0.523	0.054	0.344	0.633
$\bar{y}_{it}$ – Mean income or expenditure	246.63	90.55	55.53	671.04
	Middle East and North Africa (N=37)			
H_{it} – Headcount ($2)	0.166	0.111	0.003	0.466
G_{it} – Gini coefficient	0.380	0.042	0.301	0.474
$\bar{y}_{it}$ – Mean income or expenditure	165.26	56.59	84.02	306.33
	South Asia (N=35)			
H_{it} – Headcount ($2)	0.672	0.226	0.122	0.936
G_{it} – Gini coefficient	0.343	0.067	0.259	0.627
$\bar{y}_{it}$ – Mean income or expenditure	67.78	39.20	30.71	204.98
	Sub-Saharan Africa (N=129)			
H_{it} – Headcount ($2)	0.708	0.202	0.018	0.985
G_{it} – Gini coefficient	0.453	0.087	0.289	0.743
$\bar{y}_{it}$ – Mean income or expenditure	67.62	54.04	14.93	465.80

Notes: The table reports regional summary statistics. Mean income or expenditure is reported in 2005 PPP international dollars per month. 809 observations, 124 countries in total, unbalanced sample from 1981 to 2010.

Figure A-1 – Inequality changes and income growth, 1981–2010

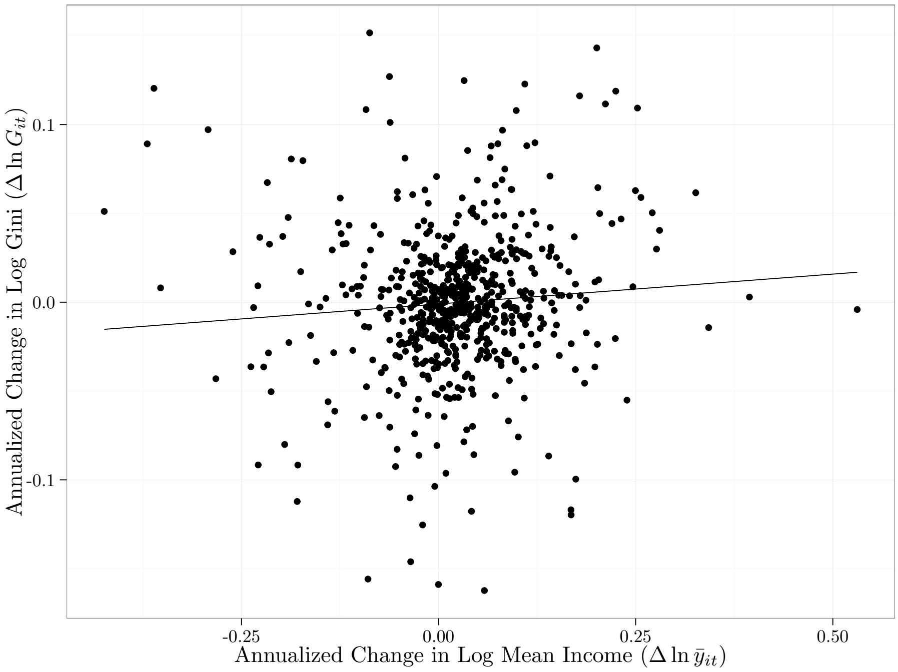

Notes: Illustration of the weak correlation between annualized changes in average income and annualized changes in inequality, as measured by proportionate changes in the Gini coefficient.

Figure A-2 – Transformed headcount ratios ($2 a day) and log Gini, by region

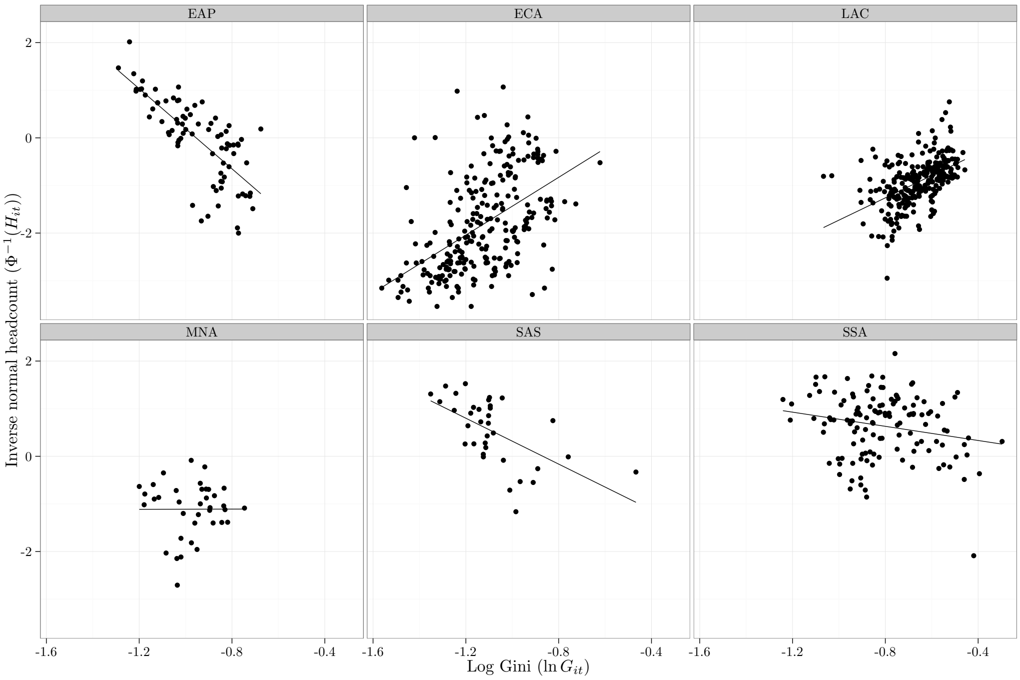

Notes: Illustration of how effectively the inverse normal cdf linearizes the conditional mean. The figure plots $\Phi^{-1}(H_{it})$ against $\ln \bar{G}_{it}$ per region, including a regression line.

Figure A-3 – Time trends in headcount ratios ($2 a day), by region

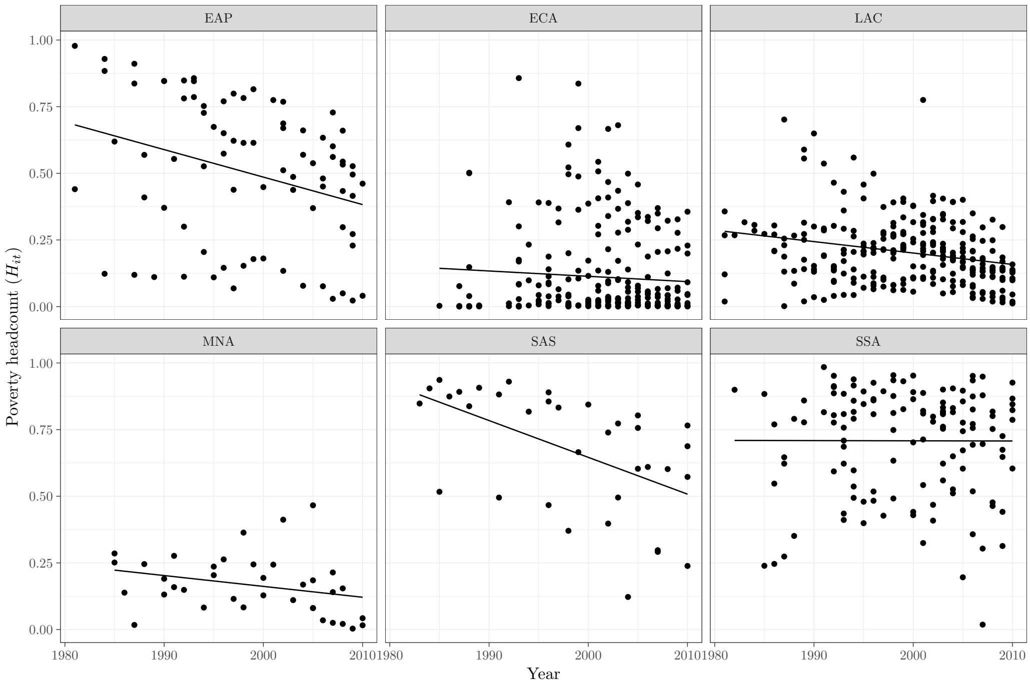

Notes: Illustration of the regional diversity in poverty levels and trends. The figure plots H_{it} against the (rounded) survey year per region, including a regression line.

A.2 Growth in personal consumption expenditures

Consumption expenditures from the national accounts (PCE) play an important role in the paper. They form the basis of the identification strategy and are used as inputs in the forecasting model. Here we highlight two major points: *(i)* the long-run trend of PCE growth in Sub-Saharan Africa has shifted upward in the mid-1990s and again in 2000; *(ii)* PCE growth has been very unsteady in the past, not only in Sub-Saharan Africa but also in all other regions apart from East Asia.

Figure A-4 illustrates the first point. It plots two series: PCE growth in the developing world without Sub-Saharan Africa, and PCE growth in Sub-Saharan Africa. The series are 10-year simple moving averages in order to smooth out cyclical variation and isolate the long-run trend. We use a formal structural change test (Andrews, 1993) to detect an unknown break point in the Sub-Saharan African trend. The test identifies the year 2000, with an associated sup-F statistic of 97.43 and an approximate p-value < 0.001. We also ran population-weighted regressions with PCE growth in all Sub-Saharan African countries as the dependent variable and a dummy variable marking all observations after 1999, plus a constant, on the right hand side. The raw data also indicates a break around the millennium. PCE growth before 2000 is estimated to be -0.82% per annum. After 2000, growth increases to 2.4% per annum – a sizable up-break of about 3.2%.

Table A-3 analyzes this temporal and regional instability of PCE growth rates in greater detail. It reports regional PCE growth rates over various time periods which are the backbone of our poverty projections. It is readily apparent why the period from 2000 to 2010 corresponds to our 'optimistic' growth scenario, 1980 to 2010 to our 'moderate' growth scenario, and 1980 to 2000 to our 'pessimistic' growth scenario. Only East Asia experiences relatively stable consumption growth throughout all sub-periods. Several other regions were plagued by slow or negative growth in the 1980s and 1990s. However, most regions exhibit at least a moderate acceleration of PCE growth in the period from 2000 to 2010 (as highlighted in Figure A-4).

Figure A-4 – Long-run trends in PCE growth – SSA versus developing world, 1985–2010

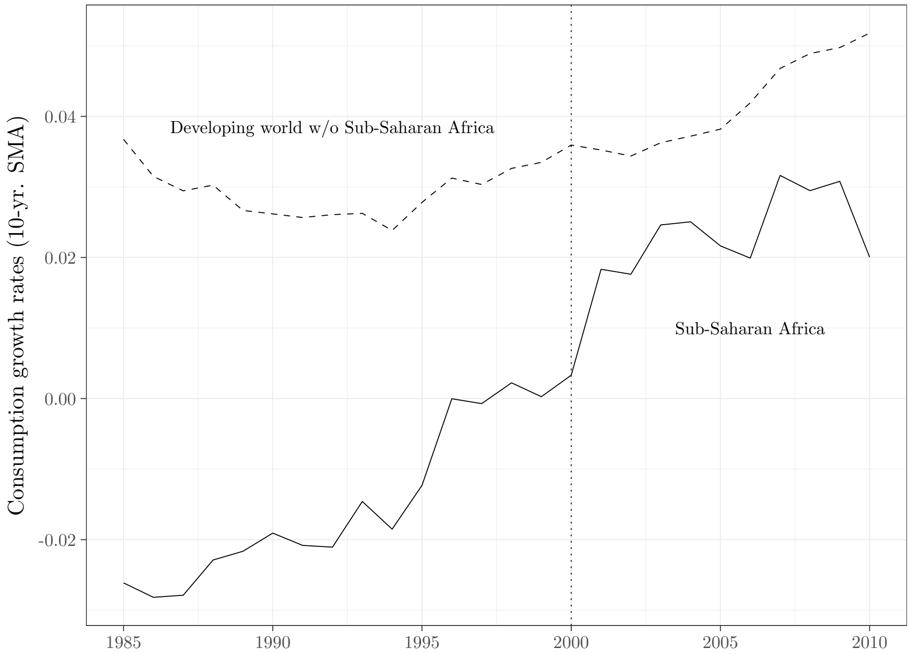

Notes: The figure plots the long-run trend in per capita consumption expenditure growth in Sub-Saharan Africa (SSA) and the developing world without SSA. The time series have been smoothed using 10-year simple moving averages (SMA). The dotted vertical line marks the year 2000 where the long run trend in SSA experienced a structural break, as indicated by a formal structural change test (Andrews, 1993).

Table A-3 – Growth in personal consumption expenditures per capita (in %), by region

	Time Period				
	2000-2010	1990-2010	1980-2010	1990-2000	1980-2000
East Asia and Pacific	5.906	5.772	5.598	5.608	5.377
	(0.813)	(0.653)	(0.725)	(0.508)	(0.677)
Europe and Central Asia	6.085	2.755	2.558	-1.225	-0.769
	(0.989)	(0.412)	(0.411)	(1.027)	(0.916)
Latin America and the Caribbean	2.444	2.219	1.445	1.931	0.677
	(0.239)	(0.140)	(0.098)	(0.337)	(0.171)
Middle East and North Africa	3.495	2.532	1.851	1.253	0.495
	(0.443)	(0.440)	(0.293)	(0.648)	(0.545)
South Asia	4.448	3.612	3.179	2.511	2.173
	(0.489)	(0.388)	(0.351)	(0.294)	(0.284)
Sub-Saharan Africa	2.382	1.419	0.698	0.016	-0.818
	(0.689)	(0.470)	(0.472)	(0.688)	(0.540)
N	123	123	123	122	122
$\bar{T}$	10.99	20.64	27.16	9.730	16.30
$N \times \bar{T}$	1352	2539	3341	1187	1989

Notes: The table reports population-weighted estimates of regional PCE growth. Cluster robust standard errors are reported in parentheses.

B Additional regression results

B.1 Strength of the first stage and fit of the second stage

In the estimation of eq. (14) in the main text we use personal consumption expenditures from the national accounts as a time-varying instrument for the observed survey means. The first-stage relationship is very strong. The Kleibergen-Paap F-statistic is 28.05 when we carry out a small sample adjustment and 31.39 without. This far exceeds even the strictest critical values for one endogenous regressor given in Stock and Yogo (2005).

Figure B-1 illustrates the strength of the first-stage relationship visually. It shows a partial regression plot of the variation in log income that is captured by personal consumption expenditures, after removing the influence of all other variables including the fixed effects. Estimates of $\widetilde{\ln \bar{y}_{it}} = \ln \bar{y}_{it} - \mathbf{x}'_{1it}\boldsymbol{\pi}_1 - \sum_{\ell=2}^{T} \delta_{T_i,\ell}(\pi_{0\ell} + \bar{\mathbf{x}}'_i\boldsymbol{\pi}_{3\ell} + \bar{\mathbf{z}}'_i\boldsymbol{\pi}_{4\ell})$ are plotted on the y-axis, and estimates of $\widetilde{\ln PCE^P_{it}} = \ln PCE^P_{it} - \mathbf{x}'_{1it}\boldsymbol{\pi}_1 - \sum_{\ell=2}^{T} \delta_{T_i,\ell}(\pi_{0\ell} + \bar{\mathbf{x}}'_i\boldsymbol{\pi}_{3\ell} + \bar{\mathbf{z}}'_i\boldsymbol{\pi}_{4\ell})$ on the x-axis. As in the main text, $\mathbf{x}_{1it}$ excludes $\ln \bar{y}_{it}$ but includes the log of Gini, while $\bar{\mathbf{x}}_i$ and $\bar{\mathbf{z}}_i$ contain the time averages of $\ln G_{it}$ and $\ln PCE^P_{it}$, respectively.

Figure B-2 shows that our fractional response framework fully captures the inherent non-linearity of the poverty data in the second estimation stage. Using our preferred specification, we plot both the observed headcount and the predicted headcount over the range of observed mean income or expenditures (left panel) and inequality (right panel). The non-linear fit is very close at either bound (near unity or near zero). The model predicts no nonsensical values, and the entire range of variation of the observed values is covered by the model predictions.

Figure B-1 – Partial regression plot – first stage

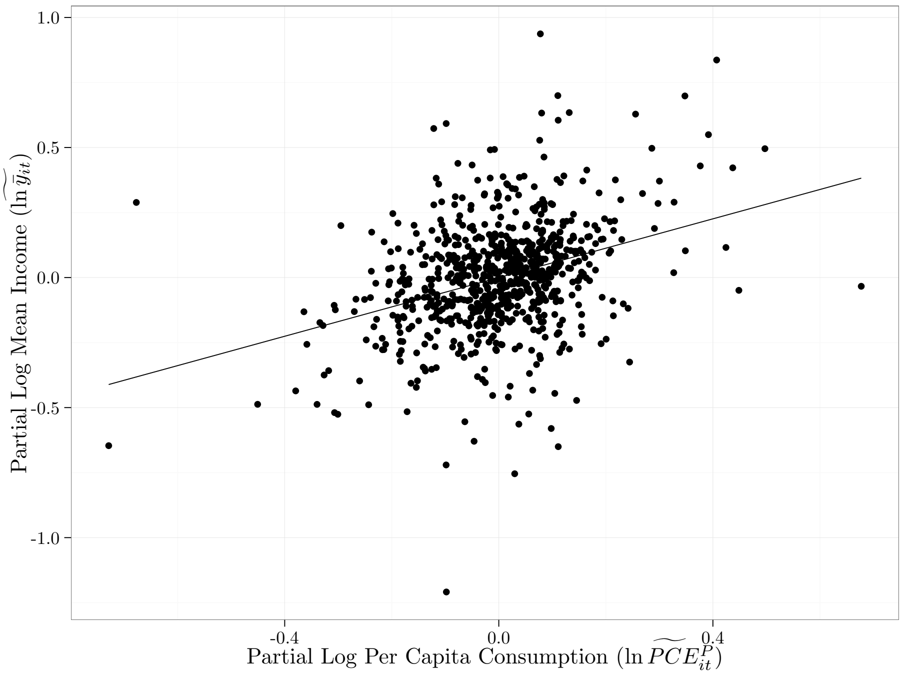

Notes: The figure plots two residual series, so that the plotted slope is identical to the slope of $\ln PCE_{it}^P$ in the first stage of the results presented in Table 2.

Figure B-2 – Data versus fitted values, preferred specification, \$2 a day

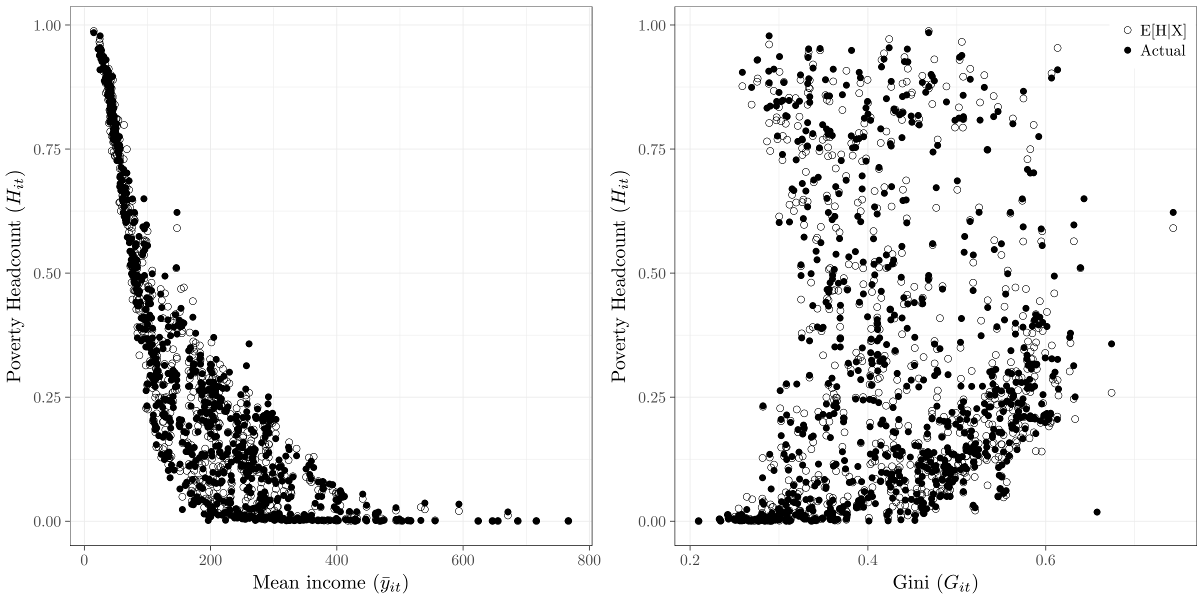

Notes: Illustration of the model fit from the estimates presented in column (3) of Table 2. The left panel shows the observed poverty rates (actual) and predicted poverty rates ($\hat{\mathrm{E}}[H|X]$) over mean incomes; the right panel plots both over the Gini coefficient.

B.2 Comparison with linear models

The fractional response model presented in this paper is very different from the standard linear approaches. Here we explicitly compare some of their features to underline the statements made in main text. As is typical for such comparisons, linearly and non-linearly estimated elasticities are very similar *on average*. The differences emerge once we use instrumental variables or examine elasticities away from the mean (which is also further explored in Online Appendix C).

Table B-1 replicates the main models as they were advanced in the literature using our data set. A key obstacle for the comparison is that these models are specified with logs or log-differences in the dependent variable, while our approach models the conditional expectation of the poverty rate directly. Hence we refrain from reporting measures of fit that are generally much lower but not comparable to ours. Column (1) uses a simple fixed effects specification with income and inequality on the right hand side. Column (2) uses annualized differences instead. Column (3) specifies the 'improved standard model' as proposed by Bourguignon (2003, see eq. (7) in the main text). Note that the coefficients on several interaction terms are not significant, suggesting that the model may not be correctly specified. Still, in all three cases, the results for the average elasticities are close to ours, although the estimated Gini elasticity tends to be a bit higher.

The next three columns use GMM techniques with several external and internal instruments to correct for measurement errors as in Kalwij and Verschoor (2007). The instrument set varies across specifications and is reported in the table notes. Column (4) repeats the simple difference specification but corrects for endogeneity. We now obtain results for the average elasticities that are very similar to those of the fractional response approach, although Hansen's J-statistic casts some doubt on the validity of the instrument set. Column (5) uses a similar instrumentation strategy for the improved standard model. Now none of the coefficients on the variables involving $\Delta \ln G_{it}$ are significant. The implied elasticities are also very different from ours. Consider a one percent increase in average income or inequality when both initial levels are at the lower quartile of the data. Column (5) suggests an income elasticity of -2.05 and a Gini elasticity of 2.68, whereas our model suggests they are -1.35 and 1.03, respectively. The linearly estimated elasticities seem to be moving in the wrong direction. Column (6) mimics the preferred specification of Kalwij and Verschoor (2007) by introducing the lagged levels on their own. Now nearly all coefficients are insignificant and the specification even implies a negative average Gini elasticity. Note that Kalwij and Verschoor (2007) do not estimate a Gini elasticity that is negative on average and most of their coefficients are highly significant. Our analysis, however, suggests that an instrument set which is often weak or not plausibly exogenous leads to two-step GMM estimates of interaction models which are very unstable and produce highly implausible estimates away from the mean.

Table B-1 – Linear models – Dependent variable: $\ln H_{it}$ or $\Delta \ln H_{it}$, \$2 a day

	OLS			Two-Step GMM		
	(1) Within R+C '97	(2) Differences R+C '97	(3) Differences Bourg. '03	(4) Differences R+C '97	(5) Differences Bourg. '03	(6) Differences K+V '07
$\Delta \ln \bar{y}_{it}$		-1.895	-0.268	-2.028	2.046	-0.362
		(0.170)	(0.617)	(0.271)	(1.043)	(3.216)
$\Delta \ln \bar{y}_{it} \times \ln(\bar{y}_{i,t-1}/z)$			-0.552		-0.995	-0.517
			(0.179)		(0.258)	(0.785)
$\Delta \ln \bar{y}_{it} \times \ln G_{i,t-1}$			1.108		3.445	2.097
			(0.671)		(1.192)	(2.315)
$\Delta \ln G_{it}$		2.336	-0.527	1.664	1.257	-8.222
		(0.311)	(1.449)	(1.008)	(4.127)	(11.185)
$\Delta \ln G_{it} \times \ln(\bar{y}_{i,t-1}/z)$			1.261		-0.315	-1.382
			(0.427)		(1.172)	(1.996)
$\Delta \ln G_{it} \times \ln G_{i,t-1}$			-1.769		-1.416	-8.164
			(1.586)		(3.929)	(8.296)
$\ln(\bar{y}_{it}/z)$	-2.114					
	(0.204)					
$\ln G_{it}$	3.024					
	(0.409)					
$\ln(\bar{y}_{i,t-1}/z)$						-0.023
						(0.037)
$\ln G_{i,t-1}$						-0.129
						(0.134)
$\bar{\varepsilon}^{H\bar{y}}$	-2.114	-1.895	-1.755	-2.028	-1.905	-2.684
$\bar{\varepsilon}^{HG}$	3.024	2.336	2.201	1.664	2.206	-2.345
$N \times \bar{T}$	648	648	648	641	641	641
N	104	104	104	102	102	102
Hansen's J (p-val.)	–	–	–	0.0418	0.579	0.639

Notes: The table reports OLS and GMM estimates of the model suggested in the previous literature. The dependent variable is the log (difference) of the poverty rate at \$2 a day (in 2005 PPPs). The panel structure is country-survey-year. All standard errors are robust to clustering at the country level. The GMM results are estimated using two-step efficient GMM. Column (4) uses as instruments ΔPCE_{it}, $PCE_{i,t-1}$, $\ln \bar{y}_{i,t-1}$ and $\ln G_{i,t-1}$. Column (5) uses as instruments ΔPCE_{it}, $PCE_{i,t-1}$, $\Delta PCE_{it} \times \ln G_{i,t-1}$, $\Delta PCE_{it} \times \ln(\bar{y}_{i,t-1}/z)$, $\ln \bar{y}_{i,t-1}$, $\ln \bar{y}_{i,t-1} \times \ln G_{i,t-1}$, $\ln \bar{y}_{i,t-1} \times \ln(\bar{y}_{i,t-1}/z)$, $\ln G_{i,t-1}$ and $\ln G_{i,t-1} \times \ln G_{i,t-1}$. Column (6) uses the same instruments as column (5) but $\ln \bar{y}_{i,t-1}$ and $\ln G_{i,t-1}$ instrument for themselves. All models include a constant (not shown) and column (1) includes a time trend (not shown). Columns (2) and (4) are similar to Ravallion and Chen (1997) but we update their approach by also including the Gini as in Adams (2004). Columns (3) and (5) are similar to the 'improved standard model 2' in Bourguignon (2003). Column (6) is in the spirit of the preferred specification in Kalwij and Verschoor (2007), who use the annualized log difference of the population size ($\Delta \ln pop_{it}$) as an additional instrument and rely on real GDP per capita instead of real per capita consumption. In our data, a first-stage F-test shows that $\Delta \ln pop_{it}$ is an extremely weak instrument. Kalwij and Verschoor (2007) also use interactions of lagged inequality and lagged income with regional dummies as instruments. However, first stage diagnostics suggest a weak IV problem (the F-stat with regional dummy interactions is always lower than without) and thus we opt for a simpler instrument set. Further, in column (5) and equation (7) we do not include the lagged levels of income and inequality. Column (6) includes them for comparison with Kalwij and Verschoor (2007).

B.3 A new test of log-normality in aggregate poverty data

The model presented in eq. (13) of the main text can be modified to nest the case that the data from each underlying household survey follow a log-normal distribution. Ignoring heterogeneity, unbalancedness and endogeneity to simplify the exposition, we may specify the conditional expectation function as

$$\mathrm{E}[H_i|\bar{y}_{it}, \sigma_{it}] = \Phi\left(\frac{\alpha + \beta \ln \bar{y}_{it} + \gamma_1 \sigma_{it}^2}{\exp(\gamma_2 \ln \sigma_{it}^2)}\right), \tag{B-1}$$

where $\ln \sigma_t^2$ is obtained using eq. (6) from the main text (which is only valid under log-normality) and enters the variance function in the denominator.

If poverty headcount ratios can be characterized by a log-normal distribution with parameters that vary over surveys (i.e., countries and years in our case), then we should find $\beta = -1$ and $\gamma_1 = \gamma_2 = 1/2$.[1] Testing these restrictions amounts to testing for log-normality in the underlying income or consumption data. To see this, simply insert these parameter values, along with $\alpha = \ln z$, to recover the right-hand side of eq. (2) from the main text. We only use the slope parameters for the joint test since we later introduce additional dummies to account for different survey types and unbalancedness. A Monte Carlo simulation suggests that such a test is accurately sized in samples such as ours.

Table B-2 applies this test to our panel of aggregated household surveys. The three specifications mimic the corresponding columns in Table 2 from the main text but introduce inequality in line with the log-normal distribution; the added variance parameter in particular diverges substantially from its log-normal value. In all three columns, we strongly reject the null of log-normality. The parameter values in column (1) are numerically close to their log-normal counterparts but begin to diverge once we account for unbalancedness and endogeneity in columns (2) and (3). The overall fit is slightly improved but the improvement is tiny once unbalancedness is accounted for. The estimated income semi-elasticity barely moves so that practical implications are limited anyway. The results strongly favor our simpler and less assumption-laden approximation. This reflects a well-known result, namely, the log-normal distribution is a good first approximation to income and poverty data but does not hold up to closer scrutiny.

[1] Gibrat's law is one way how the log-normal distribution can arise from a sequence of stochastic income shocks. Consider $\ln y_t = \ln y_{t-1} + e_t$, where e_t is a random transitory shock in log-income $\ln y_t$. As t grows the distribution of e_t determines the distribution of $\ln y_t$. Battistin, Blundell, and Lewbel (2009) recently argued that this process is better thought of in terms of permanent income and suggest that consumption is closer than income to a log-normal distribution.

Table B-2 – LN fractional probit models (QML) – Dependent variable: H_{it}, \$2 a day

	(1) Regular		(2) Unbalanced		(3) Unbalanced + Two-Step	
	H_{it}	APEs	H_{it}	APEs	H_{it}	APEs
	Mean function					
$\ln \bar{y}_{it}$	-1.059	-0.284	-1.134	-0.281	-1.206	-0.302
	(0.020)	(0.005)	(0.040)	(0.005)	(0.373)	(0.020)
σ^2_{it}	0.518	-	0.595	-	0.585	-
	(0.020)		(0.030)		(0.185)	
$\hat{\nu}_{it}$					0.065	
					(0.062)	
	Variance function					
$\ln \sigma^2_{it}$	0.430		0.372		0.380	
	(0.022)		(0.023)		(0.031)	
	Joint test of log-normality					
p-value	0.003		0.000		0.001	
	Summary statistics					
Scale Factor	0.260		0.248		0.252	
$N \times \bar{T}$	789		789		775	
N	104		104		103	
pseudo-R^2	0.997		0.998		0.998	
$\ln \mathcal{L}$	-314.9		-314.4		-312.2	
$\sqrt{MSE}$	0.0157		0.0112		0.0115	

Notes: The table reports fractional response QML estimates. The dependent variable is the poverty rate at \$2 a day (in 2005 PPPs). 20 observations with $T_i = 1$ are not used during estimation. The panel structure is country-survey-year. σ^2_{it} is the variance of the log-normal distribution which we obtained by inverting eq. (6) from the main text. The joint test of normality reports the result of a Wald test of the null hypothesis that the coefficient on $\ln \bar{y}_{it}$ is unity and the coefficients on the two σ^2_{it} terms are both one-half. In models (1) and (2), the standard errors of the coefficients are robust to clustering at the country level and the standard errors of the APEs are computed via the delta method. The standard errors of the coefficients and the APEs in model (3) account for the first stage estimation step with a panel bootstrap using 999 bootstrap replications. The linear projection in the first stage uses $\ln PCE^P_{it}$ as an instrument for $\ln \bar{y}_{it}$. The first-stage cluster-robust F-statistic in (3) is 28.05. Model (3) also excludes West Bank and Gaza entirely (2 observations) and 12 observations from ECA countries pre-1990 for lack of PCE data.

B.4 Specifications with other covariates

In policy circles it is sometimes suggested that the income and inequality elasticities can "vary" as functions of other, more substantive factors (see, e.g., Epaulard, 2003). However, if the intrinsic non-linearity arising from differences in income and in the shape of the distribution is fully accounted for, then theoretically there is no room for additional factors to play a role. Here we show that this is indeed the case when our fractional response approach is used.

Table B-3 extends our main regression results with data on institutions, human capital, access to credit and trade openness. The APEs of income and inequality are not affected by the inclusion of the additional covariates and the APEs of the latter are virtually zero. Contrary to linear approximations, the fractional response approach leaves little room for misspecification of the decomposition. In other words, the relevant link is not between some factor X and a measure of poverty, but between X and income or inequality; those are the causal relationships deserving closer theoretical and empirical attention. Thus if one is interested in the effects of, say, institutions on poverty, our recommendation is to investigate the effects of institutions on income and inequality. Distinguishing this layer of relationships is important, as the impacts of income and distributional changes themselves depend on the initial levels of income and inequality.

Table B-3 – Fractional probit models (QMLE) – Dependent variable: H_{it}, \$2 a day

	(1) Institutions		(2) Human Capital		(3) Credit		(4) Trade	
	H_{it}	APEs	H_{it}	APEs	H_{it}	APEs	H_{it}	APEs
$\ln \bar{y}_{it}$	-0.888	-0.285	-0.878	-0.284	-0.950	-0.289	-0.708	-0.302
	(0.050)	(0.012)	(0.060)	(0.011)	(0.036)	(0.009)	(0.032)	(0.012)
$\ln G_{it}$	0.779	0.250	0.805	0.261	0.765	0.233	0.581	0.248
	(0.107)	(0.028)	(0.104)	(0.027)	(0.102)	(0.027)	(0.097)	(0.033)
Executive Constraints	0.005	0.001						
	(0.005)	(0.001)						
Years of Schooling			-0.002	-0.001				
			(0.017)	(0.006)				
Private Credit / GDP					-0.007	-0.002		
					(0.040)	(0.012)		
Trade Openness							0.005	0.002
							(0.017)	(0.007)
Scale factor	0.321		0.324		0.304		0.426	
$N \times \bar{T}$	678		705		697		385	
N	85		87		93		81	
AIC	894.8		914.1		887.6		552.5	
$\ln \mathcal{L}$	-276.4		-286.1		-282.8		-163.2	
$\sqrt{MSE}$	0.0203		0.0211		0.0201		0.0233	

Notes: The table reports fractional response QML estimates. The dependent variable is the poverty rate at \$2 a day (in 2005 PPPs). The panel structure is country-survey-year. The estimation samples are reduced due to the incomplete data coverage of the covariates. Observations with $T_i = 1$ are not used in estimation. All models include time averages (CRE), time dummies, survey dummies, panel size dummies and interactions between the panel size dummies and the time averages (CRE). The time averages are recomputed for each sample size. The coefficients of the time averages of the survey dummies and time effects are constrained to be equal across the panel sample sizes. The variance function depends on the sample size. The standard errors of the coefficients are robust to clustering at the country level and the standard errors of the APEs are computed via the delta method. Data on *Executive Constraints* is from the Polity IV database. Human capital is measured as *Total Years of Schooling* from Barro and Lee (2013). We linearly interpolate the five-yearly data to an annual series. *Private Credit / GDP* measures financial development and is from Beck et al. (2010). *Trade Openness* is the *de jure* binary measure from Wacziarg and Welch (2008).

B.5 Cross-validation tests

We conduct two cross-validation tests in order to asses the out-of-sample performance of our fractional response approach. The first test focuses on regional and developing world aggregates, since these are the statistics we chose to forecast in the poverty projections. The second test focuses on how accurately we can predict the individual observations underlying these summary statistics.

Table B-4 uses the observations before 2005 as testing data to predict average regional poverty rates over the period from 2005 to 2010. The fit is remarkably close: the $\sqrt{MSE}$ for these six regional data points is 0.011 and the pseudo-R^2 is 0.998. The forecast of the developing country aggregate is off by less than 0.2 percentage points, while most regional forecasts are off by less than half a percentage point. Only in South Asia and Sub-Saharan Africa we observe forecast errors up to 2 percentage points.

Table B-5 conducts a 10-fold cross-validation where 10% of the time periods are deleted in each fold but the panel structure is otherwise kept intact. This test is considerably more demanding as we now randomly remove countries in some years, even back to 1981, and focus on how well we are able to match their actual poverty rates. Note that our panel is so heavily unbalanced (see Table A-1) that this exercise sometimes uses only one observation in a country to predict another observation which is potentially very far apart (and may use a different survey methodology). Nevertheless, we find a cross-validation error of about 3.6 percentage points and a mean absolute error (MAE) of about 2.5 percentage points – less than twice the in-sample error.

Table B-4 – Predicting the 2005–2010 period with data from before

Region (r)	Predicted $\hat{\bar{H}}_r$	True $\bar{H}_r$	Absolute Error ($\lvert\hat{\bar{H}}_r - \bar{H}_r\rvert$)
East Asia and Pacific	0.35589	0.35458	0.00131
Eastern Europe and Central Asia	0.03008	0.02612	0.00397
Latin America and Caribbean	0.12292	0.11684	0.00608
Middle East and North Africa	0.13103	0.13535	0.00431
South Asia	0.68712	0.66714	0.01998
Sub-Saharan Africa	0.68642	0.70236	0.01594
Developing world	0.32925	0.32777	0.00148

Notes: The table reports pseudo out-of-sample forecasts where the data before 2005 are used as training data to predict the observed data in the 2005 to 2010 period (testing data). The dependent variable is the poverty rate at $2 a day (in 2005 PPPs). The results have not been population-weighted and are not comparable to the 2010 baseline used in the projections.

Table B-5 – 10-fold cross-validation – preferred specification

Fold	$\sqrt{MSE}$	MAE	Pseudo-R^2	% Training	% Testing
1	0.03354	0.02649	0.991	0.962	0.038
2	0.02779	0.01958	0.992	0.892	0.108
3	0.04142	0.02814	0.985	0.911	0.089
4	0.03623	0.02765	0.988	0.865	0.135
5	0.03838	0.02471	0.986	0.855	0.145
6	0.04448	0.03260	0.984	0.932	0.068
7	0.03550	0.02464	0.988	0.900	0.100
8	0.03134	0.02427	0.989	0.871	0.129
9	0.03000	0.01957	0.989	0.907	0.093
10	0.04397	0.02533	0.982	0.904	0.096
$CV_{(10)}$ - Averages	0.03668	0.02530	0.987	0.900	0.100

Notes: The table reports the results of a 10-fold cross-validation exercise for panel data. In each fold, about 10% of the time series data are deleted while all countries remain in the sample. The remaining 90% of the data (training data) are then used to predict these 10% of observations (testing data). $\sqrt{MSE}$ is the root mean squared error, MAE is the mean absolute out-of-sample error, and the pseudo-R^2 is the squared correlation of the out-of-sample predictions with the testing data. The dependent variable is the poverty rate at $2 a day (in 2005 PPPs).

B.6 Results for the $1.25 a day poverty line

The analysis done in the paper for the $2 a day poverty line was repeated using the $1.25 a day line. The results are presented here. The main differences are in the projections.

Table B-6 reports the regression results where we progressively account for more features of the data. The results are qualitatively and quantitatively very similar to those presented in the main text. Comparing columns (1) and (2) shows that unbalancedness makes little difference in practice, but accounting for endogeneity does. The average income semi-elasticity increases from column (2) to column (3) by an amount comparable to the increase at the $2 line. The Hausman test in column (3) fails to reject the null of exogeneity at conventional levels when we bootstrap over the first stage, but rejects the null at the 5% level when we ignore the first-stage. This is still to be interpreted as evidence of endogeneity since the null hypothesis does not depend on the first stage uncertainty. The average elasticities implied by these models are larger in absolute value than in the $2 a day case ($\bar{\hat{\varepsilon}}^{H\bar{y}} \approx -2.70$ with $se = 0.286$ and $\bar{\hat{\varepsilon}}^{HG} \approx 2.56$ with $se = 0.325$), while the estimated semi-elasticities are marginally smaller in absolute value. This is to be expected since lowering the absolute poverty line is equivalent to making all countries richer. The fit is about the same at both poverty lines.

The contributions of growth and distributional change to poverty reduction are shown in Table B-7. Numerically these estimates differ from those presented in the main text but, here too, the qualitative pattern is virtually the same. Panel (a) shows that distribution played a larger role over the entire 1982 to 2010 period at the $1.25 a day poverty line than at $2 a day. About one fourth of poverty reduction was due to redistribution and three fourths were due to income growth. Once we split the sample at the turn of the millennium, this share doubles from about 15% before 2000 to more than a third in the period since 2000. Interestingly, we observe this rise in every geographical region and it occurs together with a positive correlation between the two components in every region (now including in East Asia).

Table B-8 reports the projection results for the $1.25 a day poverty line, based on the estimates of column (2) in Table B-6 and the different scenarios described in the main text. With optimistic, distribution-neutral growth, we project a global poverty rate of 9.11% in 2030, or about 655 million extremely poor people. About 76% of those live in Sub-Saharan Africa and about 17% in South Asia. Under moderate growth we project a global poverty rate of 13.2% in 2030, or about 950 million extremely poor people versus 1.2 billion in 2010. About 70% of the world's poor will live in Sub-Saharan Africa and about 23% in South Asia by 2030. Under a pessimistic, distribution-neutral growth scenario, there is almost no progress at all. In all three cases, the pace of poverty reduction slows down considerably, both in terms of relative changes and in terms of the poor population.

Table B-6 – Fractional probit models (QML) – Dependent variable: H_{it}, \$1.25 a day

	(1) Regular		(2) Unbalanced		(3) Unbalanced + Two-Step	
	H_{it}	APEs	H_{it}	APEs	H_{it}	APEs
$\ln \bar{y}_{it}$	-1.212	-0.216	-0.668	-0.218	-0.822	-0.259
	(0.056)	(0.010)	(0.038)	(0.008)	(0.186)	(0.033)
$\ln G_{it}$	1.238	0.221	0.726	0.237	0.779	0.245
	(0.121)	(0.022)	(0.074)	(0.020)	(0.177)	(0.031)
$\hat{\nu}_{it}$					0.152	
					(0.104)	
CRE (Corr. Rand. Effects)	Yes		Yes		Yes	
Survey type dummies	Yes		Yes		Yes	
Time dummies	Yes		Yes		Yes	
Panel size dummies	No		Yes		Yes	
Panel size dummies × CRE	No		Yes		Yes	
Variance function	No		Yes		Yes	
Scale factor	0.179		0.326		0.315	
$N \times \bar{T}$	768		768		754	
N	103		103		102	
pseudo R^2	0.975		0.990		0.991	
$\ln \mathcal{L}$	-172.4		-244.7		-243.5	
$\sqrt{MSE}$	0.0339		0.0214		0.0204	

Notes: The table reports fractional response QML estimates. The dependent variable is the poverty rate at \$1.25 a day (in 2005 PPPs). 21 observations with $T_i = 1$ are not used during estimation. The panel structure is country-survey-year. The \$1.25 a day sample is smaller as for 20 observation we only have data at the \$2 a day line. In columns (1) and (2), the standard errors of the coefficients are robust to clustering at the country level and the standard errors of the APEs are computed via the delta method. We include the time averages of the survey type and time dummies in columns (2) and (3), but constrain their coefficients to be equal across the panel sizes. The standard errors of the coefficients and the APEs in model (3) account for the first stage estimation step with a panel bootstrap using 999 bootstrap replications. The linear projection in the first stage uses $\ln PCE_{it}^P$ as an instrument for $\ln \bar{y}_{it}$. The first-stage cluster-robust F-statistic in column (3) is 24.40. Column (3) also excludes West Bank and Gaza entirely (2 observations) and 12 observations from ECA countries pre-1990 for lack of PCE data.

Table B-7 – Decomposition at $1.25 a day poverty line, by region

	Var(Y)	Var(D)	Cov(Y,D)	s_Y	s_D	$\sqrt{MSE}$	N
	Panel (a) Spells from 1981 to 2010						
East Asia and Pacific	2.270	0.706	0.274	72.20	27.80	0.89	12
Europe and Central Asia	1.838	0.656	0.405	67.88	32.12	0.32	38
Latin America and Caribbean	1.908	0.646	-0.780	113.43	-13.43	0.79	28
Middle East and North Africa	0.078	0.048	0.027	58.32	41.68	0.34	8
South Asia	0.686	2.606	-0.009	20.68	79.32	0.47	7
Sub-Saharan Africa	2.107	0.381	-0.246	93.21	6.79	0.47	29
All developing	2.165	0.698	0.074	74.36	25.64	0.57	122
	Panel (b) Spells from 1981 to 1999						
East Asia and Pacific	2.922	0.336	0.293	83.63	16.37	1.04	9
Europe and Central Asia	8.626	1.541	-0.305	87.07	12.93	0.50	23
Latin America and Caribbean	2.271	1.045	-0.782	84.98	15.02	0.92	26
Middle East and North Africa	0.129	0.011	-0.000	92.16	7.84	0.35	6
South Asia	0.887	0.234	0.115	74.15	25.85	0.47	4
Sub-Saharan Africa	4.775	1.821	-0.686	78.28	21.72	1.30	17
All developing	5.253	1.126	-0.269	85.32	14.68	0.90	85
	Panel (c) Spells from 2000 to 2010						
East Asia and Pacific	1.244	1.254	0.139	49.82	50.18	0.76	9
Europe and Central Asia	1.432	0.213	0.369	75.55	24.45	0.41	25
Latin America and Caribbean	0.216	0.094	0.073	63.36	36.64	0.42	19
Middle East and North Africa	0.064	0.029	0.028	61.80	38.20	0.40	5
South Asia	0.250	1.138	0.346	28.65	71.35	0.43	6
Sub-Saharan Africa	2.824	1.604	0.645	60.67	39.33	1.00	21
All developing	1.652	0.714	0.438	64.46	35.54	0.65	85

Notes: The table reports the results of the decomposition of the observed changes in the poverty rate at $1.25 a day into its growth and distribution components at the regional level. The estimates are based on column (3) of Table B-6. Panels (a) to (c) run this decomposition over different subsamples as denoted in the table. We predict the counterfactual quantities using the first and last available data for the longest runs of survey of the same type within the sample period.

Table B-8 – Projected poverty headcount ratios and poor population at $1.25 a day in 2030, by region

	Average PCE Growth								
	Optimistic (2000-2010)			Moderate (1980-2010)			Pessimistic (1980-2000)		
	Change in Inequality (Gini)								
	pro-poor	neutral	pro-rich	pro-poor	neutral	pro-rich	pro-poor	neutral	pro-rich
Panel (a) – Headcount at $1.25 a day in 2030 (in percent)									
East Asia and Pacific	0.65	0.93	1.31	0.76	1.07	1.48	0.94	1.29	1.74
Europe and Central Asia	0.12	0.16	0.21	1.21	1.45	1.71	5.17	5.74	6.44
Latin America and Caribbean	2.27	2.74	3.28	3.46	4.12	4.91	4.59	5.46	6.48
Middle East and North Africa	0.48	0.66	0.91	1.54	2.07	2.75	3.72	4.77	6.05
South Asia	4.19	5.54	7.24	8.48	10.89	13.79	12.76	15.99	19.77
Sub-Saharan Africa	32.09	35.69	39.37	43.62	47.17	50.70	51.75	55.12	58.47
Total	7.88	9.11	10.49	11.63	13.20	14.96	14.96	16.82	18.89
Panel (b) – Poor population at $1.25 a day in 2030 (in millions)									
East Asia and Pacific	14.05	20.23	28.56	16.59	23.29	32.22	20.44	28.00	37.87
Europe and Central Asia	0.59	0.76	0.97	5.72	6.86	8.11	24.47	27.16	30.45
Latin America and Caribbean	16.15	19.44	23.32	24.61	29.30	34.86	32.61	38.77	46.05
Middle East and North Africa	2.12	2.94	4.04	6.84	9.18	12.18	16.47	21.15	26.83
South Asia	83.47	110.38	144.35	169.13	217.08	275.00	254.40	318.89	394.16
Sub-Saharan Africa	449.54	499.97	551.61	611.15	660.85	710.32	725.03	772.26	819.22
Total	567.20	655.36	754.95	836.34	949.49	1076.37	1076.35	1209.95	1359.23

Notes: The table reports forecasts of the $1.25 a day poverty rate in 2030. The forecasts are based on the estimates reported in Column (2) of Table B-6 and the different growth and distribution scenarios outlined in the text. Population projections are from the World Bank's Health, Nutrition and Population Statistics database. The survey data are from the World Bank's *PovcalNet* database.

C Monte Carlo simulations

C.1 Approximation errors in the estimated elasticities

Our benchmark is a balanced panel of average incomes, Gini coefficients and poverty rates derived from a log-normal distribution. The log-normal panel DGP is defined as follows. For each replication we generate $\mu_{it} = 4.75 + e_{it}$ with $e_{it} \sim \mathcal{U}[-1.75, 1.75]$ and $\sigma_{it} = 0.8 + u_{it}$ with $u_{it} \sim \mathcal{U}[-.25, .25]$ where $i = 1, \ldots, N$ and $t = 1, \ldots, T$. The corresponding formulas for mean income, the Gini coefficient and the headcount ratio are given in Table C-1. The poverty line is $2 a day throughout.

The observed headcount is the average of a sequence of Bernoulli draws, i.e. $H_{it} \sim \mathcal{B}(B, H^*_{it})/B$, where B is the number of trials and H^*_{it} is the poverty rate drawn from the specified distribution. R is the number of Monte Carlo replications. We typically set $N = 100$, $T = 10$, $B = 1000$ and $R = 1000$. This mimics key features of the data analyzed in the main text but abstracts from measurement error and endogeneity.

The parameter values for the log-normal (and all other distributions) were chosen to match key features of the data. The simulated values approximately cover the range of observed incomes and the Gini coefficients, generate an average poverty rate around 30%, and imply an average income elasticity around two. Hence, our experiment reflects the range of observed experiences from Guinea to the Czech Republic.

To benchmark the performance of our model against the standard approach, we then compute *(i)* the average of the actual income (semi-)elasticities at each point, *(ii)* the average of the estimated (semi-)elasticities from a simplified fractional response model, and *(iii)* the estimated (semi-)elasticities from an OLS regression in differences with interactions (with and without logs on the left hand side).

Table C-2 reports the results. As expected, the log-linear and the fractional models perform similarly if we only focus on averages. Panel (a) shows that the mean relative error (MRE) of the income elasticity is around 10% in both cases. However, the fractional response model performs uniformly better once we focus on predictions away from the overall mean (median). The MRE of the fractional response model is always below or near 10% even in the tails of the distribution of $\bar{y}_{it}$, whereas the linear model with interactions is strongly biased (e.g. up to 372% at the fifth percentile of $\bar{y}_{it}$). Panel (b) focuses on the income semi-elasticity – the natural parameter of the fractional response model. Now the fractional response approach outperforms the linear model on average (with a relative error far below 1%) and remains accurate for semi-elasticities in the tails of the income distribution. The linear model produces semi-elasticities that vary very little and are heavily biased away from the mean. We only report income effects here, but the results for the Gini elasticities and semi-elasticities are comparable.

Table C-1 – Alternate distributions and DGPs

Distribution	Parameters (DGP)	Mean	Gini coefficient	Headcount ratio
		Two parameter distributions		
Log-normal	$\mu = 4.75 + \mathcal{U}(-1.75, 1.75)$ $\sigma = 0.8 + \mathcal{U}(-.25, .25)$	$\exp\left(\mu + \frac{1}{2}\sigma^2\right)$	$2\Phi\left(\frac{\sigma}{\sqrt{2}}\right) - 1$	$\Phi\left(\frac{-\ln(\bar{y}/z)}{\sigma} + \frac{1}{2}\sigma\right)$
Weibull	$a = 1.25 + \mathcal{U}(-.75, .75)$ $b = 175 + \mathcal{U}(-160, 160)$	$b\Gamma\left(1 + \frac{1}{a}\right)$	$1 - 2^{-\frac{1}{a}}$	$1 - \exp\left[-\left(\frac{z}{b}\right)^a\right]$
Fisk	see Dagum with $p = 1$			
		Three parameter distributions		
Dagum	$a = 3 + \mathcal{U}(-1.5, 1.5)$ $p = 1 + \mathcal{U}(-.5, .5)$ $b = 100 + \mathcal{U}(-75, 75)$	$\frac{b\Gamma(p + 1/a)\Gamma(1 - 1/a)}{\Gamma(p)}$	$\frac{\Gamma(p)\Gamma(2p + 1/a)}{\Gamma(2p)\Gamma(p + 1/a)}$	$\left[1 + \left(\frac{z}{b}\right)^{-a}\right]^{-p}$
Singh-Maddala	$a = 1.5 + \mathcal{U}(-.5, .5)$ $q = 4 + \mathcal{U}(-2.5, 2.5)$ $b = 350 + \mathcal{U}(-300, 300)$	$\frac{b\Gamma(1 + 1/a)\Gamma(q - 1/a)}{\Gamma(q)}$	$1 - \left[\frac{\Gamma(q)\Gamma(2q - 1/a)}{\Gamma(2q)\Gamma(q - 1/a)}\right]$	$1 - \left[1 + \left(\frac{z}{b}\right)^a\right]^{-q}$

Notes: The table summarizes the income distributions and parameters used in the Monte Carlo simulations. $\Phi(\cdot)$ is the cdf of the standard normal distribution and $\Gamma(\cdot)$ is the gamma function. Scale parameters are typically denoted by a (or μ for the log-normal). All other parameters define the shape of the distribution. The poverty line is fixed at $z = 60.83$ per month (or \$2 a day) in all simulations. The parameters of the DGPs have been chosen such that they generate an average headcount in the neighborhood of 0.30.

Table C-2 – Monte Carlo – QML versus OLS

	Actual	*Fractional response approach*			*Improved standard model*		
		Estimate	MRE	$\sqrt{MSE}$	Estimate	MRE	$\sqrt{MSE}$
		Panel (a) Income elasticity of poverty $\varepsilon^{H\bar{y}}$					
Mean	-1.9966	-1.8296	0.0837	0.1683	-1.7548	0.1211	0.2452
At percentile of $\bar{y}$							
$\ldots P_5$	-0.2869	-0.2926	0.0199	0.0063	-1.3585	3.7348	1.0726
$\ldots P_{25}$	-0.8338	-0.7773	0.0677	0.0569	-1.5340	0.8398	0.7018
$\ldots P_{50}$	-1.8945	-1.6937	0.1060	0.2015	-1.7518	0.0753	0.1563
$\ldots P_{75}$	-3.1784	-2.8122	0.1152	0.3671	-1.9698	0.3803	1.2106
$\ldots P_{95}$	-4.2812	-3.7806	0.1169	0.5017	-2.1707	0.4930	2.1116
		Panel (b) Income semi-elasticity of poverty $\eta^{H\bar{y}}$					
Mean	-0.2609	-0.2610	0.0002	0.0018	-0.2991	0.1463	0.0386
At percentile of $\bar{y}$							
$\ldots P_5$	-0.2528	-0.2528	0.0002	0.0018	-0.3625	0.4340	0.1105
$\ldots P_{25}$	-0.4962	-0.4665	0.0599	0.0298	-0.3343	0.3264	0.1624
$\ldots P_{50}$	-0.3362	-0.3587	0.0669	0.0227	-0.2991	0.1105	0.0413
$\ldots P_{75}$	-0.0844	-0.0885	0.0493	0.0045	-0.2640	2.1288	0.1797
$\ldots P_{95}$	-0.0180	-0.0127	0.2951	0.0053	-0.2358	12.1235	0.2179

Notes: The table reports the results of 1000 Monte Carlo simulations based on log-normal data with 1000 observations (100 countries over 10 years). Panels (a) and (b) show the results from applying the proposed panel QML estimator and OLS with interactions to the simulated data. Panel (a) reports elasticities, Panel (b) semi-elasticities.

C.2 Approximation errors in the estimated headcount ratios

We now show that our approach also works well when incomes are drawn from a variety of two and three-parameter distributions. We exclusively focus on the ability of our model to recover the cross-country distribution of poverty rates. The rationale behind this choice is simple. Elasticities and semi-elasticities are point-derivatives at the estimated poverty rates. If our model does a good job recovering the latter, the former are likely to be closely approximated as well. There is also a pragmatic reason. More complicated distributions often do not have closed-form solutions for the income and inequality elasticities (see Bresson, 2009, for an alternate approach using Lorenz curves). We do not compare these results to the standard linear approach, since the latter is formulated in differences and not designed to predict poverty rates directly.

We proceed in line with the simulations presented above. Table C-1 lists the DGPs and characteristics of the five chosen functional forms. The panel is still balanced with $N = 100$ and $T = 10$. We generate random data by averaging 1000 Bernoulli draws and run 1000 simulations. Table C-1 reports the corresponding results. For each of the five distributions, we compute the average poverty rate and five percentiles of the distribution. Our estimates are usually within one percentage point of the generated data, but can reach an absolute bias of 3–4 percentage points. Noticeable discrepancies only occur in the lowest and highest percentiles. There, the MRE can be sizable, although the percentage point differences remain moderate.

We also compute a standard error which focuses on the individual observations in each replication data set. Specifically, we take the square root of the average MSEs over all replications. This standard error varies from about 2 percentage points (log-normal) to 6 percentage points (Weibull), suggesting that even at the level of individual observations we are not too far off the actual data. For the log-normal and Singh-Maddala distributions, this is comparable to the $\sqrt{MSE}$ of our preferred specification in the main text.

Table C-3 – Monte Carlo – QML with various income distributions

	Estimate	Actual	Bias	MRE	$\sqrt{MSE}$
Panel (a) Log-normal DGP					
Mean of H_{it}	0.3259	0.3261	0.0001	0.0004	0.0002
$\ldots P_5$	0.0013	0.0031	0.0018	1.3999	0.0018
$\ldots P_{25}$	0.0292	0.0305	0.0013	0.0434	0.0024
$\ldots P_{50}$	0.2111	0.2140	0.0029	0.0137	0.0064
$\ldots P_{75}$	0.6136	0.6144	0.0008	0.0013	0.0070
$\ldots P_{95}$	0.8821	0.8768	-0.0053	0.0060	0.0081
Panel (b) Weibull DGP					
Mean of H_{it}	0.3465	0.3462	-0.0002	0.0007	0.0003
$\ldots P_5$	0.0654	0.0801	0.0147	0.2247	0.0148
$\ldots P_{25}$	0.1556	0.1651	0.0096	0.0616	0.0101
$\ldots P_{50}$	0.2875	0.2695	-0.0180	0.0625	0.0185
$\ldots P_{75}$	0.4669	0.4844	0.0176	0.0376	0.0194
$\ldots P_{95}$	0.8845	0.8446	-0.0399	0.0452	0.0425
Panel (c) Fisk DGP					
Mean of H_{it}	0.3089	0.3084	-0.0005	0.0015	0.0005
$\ldots P_5$	0.0254	0.0427	0.0173	0.6796	0.0174
$\ldots P_{25}$	0.0909	0.0922	0.0013	0.0145	0.0033
$\ldots P_{50}$	0.2105	0.2088	-0.0017	0.0081	0.0053
$\ldots P_{75}$	0.4817	0.4808	-0.0008	0.0017	0.0093
$\ldots P_{95}$	0.8600	0.8437	-0.0163	0.0190	0.0195
Panel (d) Dagum DGP					
Mean of H_{it}	0.3249	0.3239	-0.0010	0.0030	0.0010
$\ldots P_5$	0.0173	0.0354	0.0181	1.0511	0.0182
$\ldots P_{25}$	0.0940	0.1002	0.0061	0.0649	0.0069
$\ldots P_{50}$	0.2407	0.2346	-0.0061	0.0254	0.0079
$\ldots P_{75}$	0.5120	0.5067	-0.0052	0.0102	0.0105
$\ldots P_{95}$	0.8720	0.8559	-0.0162	0.0185	0.0192
Panel (e) Singh-Maddala DGP					
Mean of H_{it}	0.3416	0.3423	0.0006	0.0019	0.0006
$\ldots P_5$	0.0548	0.0519	-0.0029	0.0529	0.0036
$\ldots P_{25}$	0.1380	0.1354	-0.0027	0.0194	0.0038
$\ldots P_{50}$	0.2689	0.2749	0.0060	0.0221	0.0069
$\ldots P_{75}$	0.4904	0.5061	0.0157	0.0319	0.0167
$\ldots P_{95}$	0.8733	0.8396	-0.0337	0.0386	0.0352

Notes: The table reports the results of 1000 Monte Carlo simulations based on the DGPs given in Table C-1 with 1000 observations (100 countries over 10 years). Panels (a) to (e) show results obtained by applying the proposed panel QML estimator under each specified income distribution.

References

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

Andrews, D. W. K. (1993). Tests for parameter instability and structural change with unknown change point. *Econometrica 61*, 821–856.

Barro, R. J. and J. W. Lee (2013). A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics 104*, 184–198.

Battistin, E., R. Blundell, and A. Lewbel (2009). Why is consumption more log normal than income? Gibrat's law revisited. *Journal of Political Economy 117*(6), 1140–1154.

Beck, T., A. Demirgüç-Kunt, and R. Levine (2010). Financial institutions and markets across countries and over time: The updated financial development and structure database. *World Bank Economic Review 24*(1), 77–92.

Bourguignon, F. (2003). The growth elasticity of poverty reduction: Explaining heterogeneity across countries and time periods. In T. S. Eicher and S. J. Turnovsky (Eds.), *Inequality and Growth: Theory and Policy Implications*, pp. 3–26. Cambridge, MA: MIT Press.

Bresson, F. (2009). On the estimation of growth and inequality elasticities of poverty with grouped data. *Review of Income and Wealth 55*(2), 266–302.

Epaulard, A. (2003). Macroeconomic performance and poverty reduction. Working Paper 03/72, International Monetary Fund.

Kalwij, A. and A. Verschoor (2007). Not by growth alone: The role of the distribution of income in regional diversity in poverty reduction. *European Economic Review 51*(4), 805–829.

Ravallion, M. and S. Chen (1997). What can new survey data tell us about recent changes in distribution and poverty? *World Bank Economic Review 11*(2), 357–382.

Stock, J. H. and M. Yogo (2005). Testing for weak instruments in linear IV regression. In D. W. Andrews and J. H. Stock (Eds.), *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Chapter 5, pp. 80–108. Cambridge University Press.

Wacziarg, R. and K. H. Welch (2008). Trade liberalization and growth: New evidence. *World Bank Economic Review 22*(2), 187–231.