

Grimm, Michael

Working Paper

Removing the anonymity axiom in assessing pro-poor growth

IAI Discussion Papers, No. 113

Provided in Cooperation with:

Ibero-America Institute for Economic Research, University of Goettingen

Suggested Citation: Grimm, Michael (2005) : Removing the anonymity axiom in assessing pro-poor growth, IAI Discussion Papers, No. 113, Georg-August-Universität Göttingen, Ibero-America Institute for Economic Research (IAI), Göttingen

This Version is available at:

<https://hdl.handle.net/10419/27402>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**Ibero-Amerika Institut für Wirtschaftsforschung
Instituto Ibero-Americano de Investigaciones Económicas
Ibero-America Institute for Economic Research
(IAI)**

**Georg-August-Universität Göttingen
(founded in 1737)**



Diskussionsbeiträge · Documentos de Trabajo · Discussion Papers

Nr. 113

**Removing the anonymity axiom
in assessing pro-poor growth**

Michael Grimm

July 2005

Removing the anonymity axiom in assessing pro-poor growth

With an application to Indonesia and Peru

Michael Grimm*

Department of Economics, University of Göttingen
DIAL Paris and DIW Berlin

First draft: February 24; This draft: June 13, 2005

Abstract

The recent focus on ‘pro-poor growth’ led also to an intense debate on how exactly to define and to measure pro-poor growth. However, all suggested measures, irrespective whether they use the absolute or the relative definition of pro-poor growth have in common that they are based on the anonymity axiom, i.e. they do not distinguish between changes in horizontal and vertical inequality. That means usual assessments of pro-poor growth look at distributional changes over time and ignore how specific groups or households moved. Such a perspective may provide a very incomplete picture given that the common objective of most studies investigating the pro-poor growth is to test whether specific policy reforms were beneficial to the initially poor or not. Using panel data from Indonesia and Peru, this paper analyzes and illustrates empirically the implications of removing the anonymity axiom from measurements of pro-poor growth. It is shown that postulating anonymity, when assessing pro-poor growth can lead to misleading conclusions on how a specific policy affected the incomes of the initially poor. For both countries, the analysis shows substantial convergence to the mean, which is, at least for the case of Indonesia, robust to measurement error in the expenditure data.

JEL Classification: D31, D63, O15.

Key words: Anonymity axiom, pro-poor growth, income mobility, horizontal equity, inequality, decomposition.

*University of Göttingen, Department of Economics, Platz der Göttinger Sieben 3, 37073 Göttingen, Germany; Email: mgrimm@uni-goettingen.de.

I thank Javier Herrera and Sandrine Mesplé-Soms for having made available their Peruvian household data. Moreover, this paper benefited from fruitful discussions with Denis Cogneau and Stephan Klasen. Of course, any remaining errors and omissions remain my own responsibility.

1 Introduction

The recent focus on ‘pro-poor growth’ in development economics and politics led also to an intense debate on how exactly to define and to measure pro-poor growth (see e.g. Duclos and Wodon 2004; Klasen 2004; Ravallion 2004a). A key point in this debate is whether pro-poor growth should be defined in ‘absolute’ or in ‘relative’ terms of poverty reduction. According to the absolute definition growth is considered as being pro-poor whenever the incomes of the poor increase. In contrast, the relative definition requires that the growth rate of income is higher among the poor than among the non-poor, i.e. inequality must decrease. However, all suggested measures, irrespective whether they use the absolute or the relative definition of pro-poor growth have in common that they are based on the anonymity axiom, i.e. they do not distinguish between changes in horizontal and vertical poverty and inequality. ‘Horizontal inequality’ refers to inequality between initially ‘equals’, i.e. individuals having had initially the same income (Plotnick 1982). ‘Vertical inequality’ refers to inequality between initially ‘unequals’, i.e. inequality between different income groups.

More precisely, usual assessments of pro-poor growth look at distributional changes over time and ignore how specific groups or households moved. In other words, two distributions are treated as equally good if, after income is redistributed among households, the overall distribution is the same. The ways of how the observed distributions were obtained are deemed irrelevant. However, such a perspective may provide a very incomplete picture. The common objective of most studies investigating the pro-poorness of growth is to test whether specific policy reforms were beneficial to the poor or not. More generally, to evaluate the effectiveness of reforms one would like to know which groups benefited or lost and how much. Likewise, one would like to know, if individuals under the poverty line before and after the reform are roughly the same and thus poverty is a rather chronic state, or, in contrast, if mobility is high and poverty is rather a transient phenomenon. Issues of chronic poverty and income mobility have recently received considerable attention (e.g. Fields and Ok 1996; Hulme and Shepherd 2003), however they have so far not been considered in the framework of pro-poor growth. An exception is the study by Jenkins and Van Kerm (2003), but they analyze the issue for the USA and Germany, and not in the context of poverty reduction in the developing world.

The following example shows that these issues are of particular importance when assessing pro-poor growth. Take the simple case, where an income distribution observed in t can be divided into two equal sized groups: the ‘poor’ and the ‘rich’. Let us further assume that between t and $t + 1$ the poor see their incomes increase to a level which is above the level of the initially rich in t and the rich see their incomes decrease to a level which corresponds exactly to the level of the initially poor in t . Looking

only at marginal distributions we would judge such a growth pattern as not pro-poor, both, according to the absolute and the relative definition. However, looking at the group-specific trajectories, this growth pattern could be judged as being clearly pro-poor. This very simple example illustrates that postulating anonymity, when assessing pro-poor growth may result misleading conclusions on how a specific policy affected the incomes of the initially poor. However, obviously a clear-cut answer whether such a growth process can be called pro-poor or not cannot be given. It depends on the value judgments one might want to accept. For instance, utilitarianism and the Pareto principle may justify the unequal treatment of equals, but we may find it unfair that following a reform people at similar initial incomes are rewarded very differently. Ravallion (2004b) pointed out that ‘anti-globalizers’ seem to focus more on the losers amongst the poor and those vulnerable to poverty and therefore on horizontal inequality, whereas ‘globalizers’ focus more on aggregate inequality explaining why both groups may conclude so differently on the distributional consequences of international trade. Moreover, the *Millennium Development Goal One*, which requires to halve poverty by half before 2015, clearly focuses on aggregate poverty.

The objective of this paper is to analyze and illustrate empirically the implications of removing the anonymity axiom from measurements of pro-poor growth. Given that the empirical distinction between horizontal and vertical shifts in the income distribution requires not only information on the marginal distributions of income under alternative policies, but also on the joint distributions of income across these policy states, this paper will also offer a brief discussion of how such joint distributions can be recovered when panel data is not available.

The remainder of the paper is organized as follows. Section 2 discusses some usual measurements of pro-poor growth and analyzes the implications of postulating and removing the anonymity axiom. This section also offers various decompositions, which can help to understand more deeply distributional changes in favor of the poor. Section 3 illustrates these problems using panel data for Indonesia and Peru. Section 4 concludes and discusses briefly some methods for policy analysis allowing to remove the anonymity axiom when panel data is not available.

2 The anonymity axiom and measurements of pro-poor growth

2.1 The anonymity axiom

Assuming an income distribution over n individuals enjoying each an income y_i , anonymity postulates that all permutations of personal labels are

regarded as distributionally equivalent, i.e.:

$$(y_1, y_2, y_3, \dots, y_n) \sim_I (y_2, y_1, y_3, \dots, y_n) \sim_I (y_1, y_3, y_2, \dots, y_n).$$

This axiom, sometimes also called ‘symmetry’, requires that the underlying social welfare function uses only the information about the income variable and not about, for example, some other characteristics which might be discernible in a sample or an enumeration of the population (Cowell 2000). This assumption is usually invoked for welfare orderings, whether we look at inequality (Atkinson 1970) or at poverty (Greer, Foster and Thorbecke 1984). However, this axiom is neither trivial nor self-evident, and for certain purposes it could make sense to remove it.

2.2 Measurements of pro-poor growth

Various measures have been suggested to measure pro-poor growth.¹ I pick up two of them: The ‘growth incidence curve’ and the ‘rate of pro-poor growth’. Both were suggested by Ravallion and Chen (2003) and are widely used in the empirical literature. They are relatively intuitive and therefore convenient for illustrative purposes. First I analyze these measures for the case where the anonymity axiom is postulated. Afterwards, I study their properties when the anonymity axiom is removed.

2.2.1 Postulating anonymity

When comparing two income distributions observed in $t-1$ and t , the growth rate in income of the p th quantile, $g_t(p)$ can be written as:

$$g_t(p) = \frac{y_t(p)}{y_{t-1}(p)} - 1. \quad (1)$$

Letting p vary from p_1 to p_{\max} , $g_t(p)$ traces out what Ravallion and Chen (2003) called the ‘growth incidence curve’ (GIC).² Denoting γ_t the growth rate in mean income, it is evident from Equation (1) that if the Lorenz curve does not change then, $g_t(p) = \gamma_t$ for all p . Conversely, $g_t(p) > \gamma_t$ if and only if the ratio of the income at p , $y(p)$, and the mean income increases between $t-1$ and t . If $g_t(p)$ is a decreasing (increasing) function for all p then inequality falls (rises) over time for all inequality measures satisfying the Pigou-Dalton transfer principle. If the GIC lies above zero everywhere ($g_t(p) > 0 \forall p$) then there is first-order dominance of the distribution at date

¹See, for instance McCulloch and Baulch (2000), Kakwani, Khandker and Son (2003), Kakwani and Pernia (2000), Ravallion and Chen (2003) and Klasen (2004).

²The growth rate in income of the p th quantile can equivalently be written using the slopes of the Lorenz curves $L'(p)$ observed in t and $t-1$ as well as the corresponding growth rate in mean income γ : $g_t(p) = \frac{L'_t(p)}{L'_{t-1}(p)}(\gamma_t + 1) - 1$ (Ravallion and Chen 2003).

t over $t-1$. If the GIC switches sign then one cannot in general infer whether higher-order dominance holds by looking at the GIC alone (Ravallion and Chen 2003).

Using the concept of the GIC, Ravallion and Chen (2003) define the rate of pro-poor growth (*RPPG*) as the area under the GIC up to the headcount index, H , which gives the proportion of all individuals having an income below or equal to the poverty line, z :³

$$RPPG = \frac{1}{H_{t-1}P} \sum_{p=1}^{PH_{t-1}} g_t(p), \quad (2)$$

where P stands for the total number of quantiles. It can be shown that the *RPPG* corresponds to a change over time in the Watts poverty index, W , i.e. $RPPG = -\Delta W$.⁴ It is important to note that the *RPPG* is derived from the mean of the growth rates at all percentiles up to the headcount index, which is not the same as the growth rate of mean income of the poor. *RPPG* collapses to the growth rate in the overall mean (γ_t) if all incomes grow at the same rate. In this case inequality remains constant. Conversely if the mean of the growth rates at all percentiles exceeds the mean growth rate, inequality decreases, in the opposite case, inequality increases. Besides the anonymity axiom *RPPG* satisfies the focus axiom (the measure is invariant to income changes for the non-poor), the population invariance axiom (adding a replication of a population to that same population has no impact), the transfer axiom (inequality-reducing transfers amongst the poor are poverty reducing), but not, as shown by Kakwani and Son (2002), the monotonicity axiom (any income loss to the poor increases poverty), because *RPPG* is measured by integrating up to the headcount index in the initial period, i.e. in case poverty rose (declined) *RPPG* does not take into account those positions, which increased (decreased) the headcount index between $t-1$ and t .

2.2.2 Removing anonymity

So far it was (implicitly) assumed that we observe one income distribution in $t-1$, ($F(y_{i,t-1})$) and one in t , ($F(y_{j,t})$), where i and j do not refer necessarily to the same individuals or where at least no information is available to follow individuals over time. Now, I assume explicitly that this information is available and that it is possible to infer the joint income distribution $F(y_{i,t-1}, y_{i,t})$ for a fixed population, i.e. individuals cannot only be ordered

³Throughout the analysis I assume that there is no ambiguity about the poverty line. It is defined in absolute terms and remains constant in real terms over time.

⁴Where $W_t = 1/P \sum_{p=1}^{PH_t} \log[z/y_t(p)]$, or, in words, the Watts Index is the population mean of the log of the ratio of the poverty line to censored income, where the latter is the actual income for those below the poverty line and the poverty line for those above it.

by their income level y , but also according to some other personal circumstances revealing their identity or membership to group Ω_h , where h is a criteria classifying individuals into up to $i = 1, \dots, N$ groups. For instance, suppose we can order individuals, observed in $t - 1$ and t , according to the group membership $\Omega_{p(y_{t-1})}$ defined by the income quantile $p(y_{t-1})$ they belonged to in $t - 1$. This information allows to order individuals in ascending order according to their initial income quantile $p(y_{t-1})$ and to compute the quantile specific mean incomes and growth rates in income where each quantile comprises the same individuals in $t - 1$ and t :

$$g_t(p(y_{t-1})) = \frac{y_t(p(y_{t-1}))}{y_{t-1}(p(y_{t-1}))} - 1. \quad (3)$$

As before, letting p vary from p_1 to p_{\max} , $g_t(p(y_{t-1}))$ traces out a GIC. To distinguish this GIC from the one defined by Ravallion and Chen (2003), I denote it in what follows ‘IGIC’, for ‘Individual Growth Incidence Curve’. As for the GIC, the IGIC is a horizontal line if $g_t(p(y_{t-1})) = \gamma_t$ for all $p(y_{t-1})$, i.e. the individuals in each quantile see their incomes grow with the average growth rate. If $g_t(p(y_{t-1})) > 0$ ($g_t(p(y_{t-1})) < 0$) for all $p(y_{t-1})$, then each group is richer (poorer) in t than in $t - 1$. Conversely, $g_t(p(y_{t-1})) > \gamma_t$ if and only if the ratio of the income at $p(y_{t-1})$ and the mean income increases between $t - 1$ and t . However using the concept of the IGIC it is not true anymore that if $g_t(p(y_{t-1}))$ is a decreasing (increasing) function for all $p(y_{t-1})$ then inequality falls (rises) over time for all inequality measures satisfying the Pigou-Dalton transfer principle. This is because individuals in t are not anymore ordered in ascending order of their income, i.e. going along the quantiles $p(y_{t-1})$ is not going along richer and richer individuals in t . It might be that the initially rich end up poorer as the initially poor and the initially poor end up richer as the initially rich. In this case the IGIC would have an decreasing slope and the GIC a positive slope, i.e. inequality would increase. The difference is that the GIC compares two distributions quantile by quantile, whereas the IGIC reflects the transition between the distributions observed in $t - 1$ and t , i.e. income growth *and* income mobility. To evaluate such a change in the income distribution, one might apply some kind of compensation criteria (e.g. Hicks 1939 or Kaldor 1939), i.e. such a change is desirable if the ‘new rich’ could compensate the ‘new poor’ in a way that every group $\Omega_{p(y_{t-1})}$ is as least well off in t than $t - 1$ and at least on group is better off in t than $t - 1$. Things become really difficult, when the initially poor end up slightly poorer as the initially rich and the initially rich poorer than the initially poor. Then even applying a kind of compensation principle would mean—equal group sizes assumed—we have a welfare loss. Obviously such considerations will quickly bring us back to the usual GIC or, conversely, call for a more specific social welfare function. In Section 2.3 this issue will be analyzed in more detail using decomposition techniques.

To compute the *RPPG* for the ICIC, *IRPPG* in what follows, we may

integrate the area under the IGIC up to the headcount index of the initial period, H_{t-1} . That means we integrate the growth of income for all those individuals who had an income below or equal to the poverty line z in $t-1$. More precisely, integrating over the IGIC, implies to integrate over the same individuals in $t-1$ and t , independent whether they have still an income below the poverty line z or not in t . It follows, that the *IRPPG* can be written as

$$IRPPG = \frac{1}{H_{t-1}P} \sum_{p_{t-1}=1}^{p_{H_{t-1}}} g_t(p_{t-1}). \quad (4)$$

Obviously, we may have individuals who had an income above z in $t-1$, but who have one below z in t . These individuals would not enter the *IRPPG* (similar to the failure of the *RPPG* to satisfy the monotonicity axiom). Hence, computing *IRPPG* for the IGIC implies to focus on those initially poor. This can be taken as a special variant of the above mentioned ‘focus axiom’ and might be justified on Rawlsian grounds (Rawls 1971).

2.3 Decomposition of pro-poor growth

To understand the mobility process which separates the GIC and IGIC and which is hidden if anonymity is postulated, changes in income distribution and poverty can be decomposed in components measuring somehow general income growth and components measuring somehow mobility across the income distribution. This might also help to make value judgements about specific distributional changes. I use three approaches here. The first one follows Ravallion and Lokshin (2004) and consists in decomposing changes in the mean log deviation into changes in horizontal and vertical inequality. The second one follows Jenkins and Van Kerm (2003) and consists in decomposing any measure of the generalized Gini class of indices in a component summarizing mobility in the form of reranking and a component summarizing progressivity in income growth. The third one is—to my knowledge—new and consists in decomposing changes in poverty measured by the FGT class of indices or the Watts index into components summarizing up-ward mobility of the initially poor, income growth among the initially poor and down-ward mobility of the initially non-poor.

2.3.1 Decomposition of inequality changes into horizontal and vertical components

One can find an intense debate in the literature on how to define and to measure horizontal equity (see e.g. Auerbach and Hassett 2002; Duclos and Lambert 2000; Jenkins and Lambert 1999; Kaplow 1989; Plotnick 1982). The issue is in particular discussed in the literature on taxation. Among other things, this debate turns around the question what one means by ‘equals’ and ‘treating equals equally’. Ravallion and Lokshin (2004) define

as ‘equals’ those individuals with an initial equal income and regard two equal individuals as treated equally if they benefited both in the same way of growth and redistribution, i.e. if their gain conditional on their initial income is the same. More precisely, averaging across all individuals i the absolute change in income, $b_i = y_{it} - y_{it-1}$, one can calculate the conditional mean impact given by:

$$b_i^c = E_i(b_i|y = y_{it-1}). \quad (5)$$

Deviations of the expected impact conditional on initial income create horizontal inequality.

Using this definition any change in a decomposable inequality measure can be decomposed in a vertical and horizontal component. Taking the mean log deviation $MLD = 1/n \sum_{i=1}^n \ln(\bar{y}/y_i)$, the following decomposition can be made:

$$\begin{aligned} MLD_t - MLD_{t-1} &= \frac{1}{n} \sum_{i=1}^n \ln \left(\frac{1 + \bar{b}/\bar{y}_{t-1}}{1 + b_i/y_{it-1}} \right) \\ &= \frac{1}{n} \sum_{i=1}^n \ln \left(\frac{1 + \bar{b}/\bar{y}_{t-1}}{1 + b_i^c/y_{it-1}} \right) + \frac{1}{n} \sum_{i=1}^n \ln \left(\frac{1 + b_i^c/y_{it-1}}{1 + b_i/y_{it-1}} \right) \end{aligned} \quad (6)$$

where the first term can be interpreted as the vertical component and the second term as the horizontal component of the observed change in inequality. If ‘equals’, i.e. all individuals with the same initial income, receive the same benefit $b_i = b_i^c$ the horizontal component is zero. Conversely, when the relative expected benefit for each household b_i^c/y_{it-1} corresponds exactly to the relative increase in the mean income \bar{b}/\bar{y}_{t-1} the vertical component is zero.

To empirical implementation of this approach requires to estimate the expected impact $E_i(b_i|y = y_{it-1})$ (Equation 5) for each household. This will be done—again following Ravallion and Lokshin (2004)—using a non-parametric local regression method suggested by Cleveland (1979).

One may argue that period-income is only a very bad criteria to define ‘equals’ to measure horizontal inequality and it might be also interesting to perform this decompositions by groups defined according to their permanent income, wealth or socio-economic status.

2.3.2 Decomposition of inequality changes into progressivity of income growth and reranking

The second decomposition approach I use is based on the idea, that changes in the income distribution over time can be additively decomposed into terms representing the progressivity of income growth (P), i.e. whether income is pro-poor rather than pro-rich and the extent of reranking (R). Jenkins and

Van Kerm (2003) decompose in this way the Gini coefficient, but they show that this can be done for any member of the generalized Gini index:

$$G_t - G_{t-1} = R - P \quad (7)$$

where

$$G_t = 1 - \frac{2}{N^2} \sum_{i=1}^n (N - i_t + 1) \frac{y_{it}}{\bar{y}_t},$$

$$R = \frac{2}{N^2} \sum_{i=1}^n (i_t - i_{t-1}) \frac{y_{it}}{\bar{y}_t}$$

and

$$P = \frac{2}{N^2} \sum_{i=1}^n (N - i_{t-1} + 1) \left(\frac{y_{it}}{\bar{y}_t} - \frac{y_{it-1}}{\bar{y}_{t-1}} \right)$$

To compute G , R and P the N individuals are ordered in ascending order according to their initial income y_{it-1} . The indices i_t and i_{t-1} stand not only for the index of summation, but also for the rank of each household in the income distribution observed in $t - 1$ and t respectively.

Remember that the Gini coefficient is a weighted average of each individual's relative income, y_{it}/\bar{y}_t , where the weight is given by the reversed rank in the income distribution ($N - i_t + 1$). Therefore R is a relative-income weighted average of changes in ranks of individuals across the income distribution. Or, put differently, it can be interpreted as an index of mobility in the form of reranking. Clearly, when there is no reranking $R = 0$. By contrast R takes its maximum value equal to $2 \times G_t$ when income ranks are totally reversed, so that the poorest household in year $t - 1$ is the richest household in year t , the second poorest becomes the second richest, and so on (Jenkins and Van Kerm 2003). Given the focus on ranks as social weights, it is clear that this decomposition can only be done for this specific class of inequality measures.

P is a rank or 'social-weighted' average of the changes in relative income between years $t - 1$ and t and summarizes the progressivity of income growth across the base year income distribution.⁵ When everyone experiences equiproportionate income growth, relative incomes remain constant, and $P = 0$. If $P > 0$ income growth is concentrated more among the poorer individuals than the richer individuals, i.e. inequality tends to decrease and growth is pro-poor in the relative sense. By contrast, if $P < 0$ income growth is concentrated more among the richer individuals than the poorer individuals, i.e. inequality tends to increase and growth is not pro-poor in the relative sense (but might be in the absolute sense). When aggregate income growth is negative, $\gamma < 0$, then income growth is pro-poor in the relative

⁵The decomposition could also be written in a form using the ranks (social weights) of the final income distribution.

sense (not in the absolute sense) if the income losses are concentrated more among the richer individuals. P is bounded by $G_{t-1} - 1$ (when the richest household in $t - 1$ obtains all the income in year t) and $G_{t-1} + 1$ (when the poorest household in year $t - 1$ obtains all the income in year t). Hence Equation (7) states that inequality is reduced by progressive income growth unless more than offset by concomitant income mobility (Jenkins and Van Kerm 2003).

2.3.3 Decomposition of poverty changes into income growth, upward and downward mobility

The third decomposition approach I use consists in decomposing changes in poverty measured by the FGT class of indices (Foster, Greer and Thorbecke 1984) into components summarizing up-ward mobility of the initially poor, income growth among the initially poor and down-ward mobility of the initially non-poor. Hence, taking the FGT poverty indicators and defining the following set of 0/1 indicator variables:

$$\begin{aligned}\xi_i &= 1 \text{ if the individual } i \text{ was poor in } t - 1 \text{ and } t \text{ (} \textit{stayer}), \\ \pi_i &= 1 \text{ if the individual } i \text{ was poor in } t - 1 \text{ and non-poor } t \text{ (} \textit{mover}), \\ \psi_i &= 1 \text{ if the individual } i \text{ was non-poor in } t - 1 \text{ and poor } t \text{ (} \textit{joiner}),\end{aligned}$$

we can write the change in poverty between $t - 1$ and t as:

$$\begin{aligned} & P_{\alpha,t} - P_{\alpha,t-1} \\ &= \frac{1}{n} \left[\sum_{i=1}^n \left(\left(\frac{z - y_{it}}{z} \right)^\alpha \max(\xi_i, \psi_i) - \left(\frac{z - y_{it-1}}{z} \right)^\alpha \max(\xi_i, \pi_i) \right) \right], \quad (8)\end{aligned}$$

with $\alpha > 0$ and z the in real terms time-invariant poverty line. $\alpha = 0$ yields the headcount index, i.e. the proportion of poor individuals, and $\alpha = 1$ the poverty gap ratio, i.e. the average distance of the poor to the poverty line, where for the non-poor this distance is set to zero.

From this we can derive the following decomposition:

$$\begin{aligned} & P_{\alpha,t} - P_{\alpha,t-1} \\ &= \frac{1}{n} \left[\sum_{i=1}^n \left(\left(\frac{z - y_{it-1}}{z} \right)^\alpha \xi_i - \left(\frac{z - y_{it-1}}{z} \right)^\alpha \max(\xi_i, \pi_i) \right) \right] \\ &= \frac{1}{n} \left[\sum_{i=1}^n \left(\left(\frac{z - y_{it}}{z} \right)^\alpha - \left(\frac{z - y_{it-1}}{z} \right)^\alpha \right) \xi_i \right] \\ &\quad + \frac{1}{n} \sum_{i=1}^n \left(\frac{z - y_{it}}{z} \right)^\alpha \psi_i, \quad (9)\end{aligned}$$

where the first component gives the change in poverty, which is due to the up-ward mobility of those individuals who were poor in $t - 1$ and non-poor

in t (*movers*) while for those remaining under the poverty line (*stayers*) income is kept at the initial level of $t - 1$. The second component gives the change in poverty, which is due to changes in income among the *stayers* (or chronic poor). The third component gives the change in poverty, which is due to the down-ward mobility and income contraction of individuals who were initially non-poor (*joiners*). If the headcount index (FGT0), i.e. $\alpha = 0$, is retained as poverty indicator, the second component is of course zero.

An completely equivalent decomposition can be performed with the Watts index, which has a direct link to the GIC, because its negative change corresponds to the *RPPG* (see Section 2.2.1).

3 An empirical illustration for Indonesia and Peru

3.1 Data

To illustrate the implications of removing the anonymity axiom of measurements of pro-poor growth, I use longitudinal data for Indonesia and Peru.

For Indonesia, I use all three existing waves of the Indonesian Family Life Survey conducted by RAND, UCLA and the University of Indonesia's Demographic Institute in 1993 (IFLS1), 1997 (IFLS2) and 2000 (IFLS3). The IFLS is an ongoing longitudinal socioeconomic and health survey. It is representative of 83% of the Indonesian population living in 13 of the nation's current 26 provinces. The IFLS is judged as having a very high quality, among other things, because individuals who moved are tracked to their new location and, where possible, interviewed there. Hence, this procedure ensured that the re-contact rate in the IFLS3 was 95.3% of IFLS1 households. Hence, nearly 91% of IFLS1 households are complete panel households.⁶ Using the three waves, I built two panels, one from 1993 to 1997 (6,723 households; 31,324 individuals) and one from 1997 to 2000 (7,187 households; 32,314 individuals).⁷ I use real household expenditure per capita as the welfare measure, or income measure in what follows. Expenditure is expressed in 1993 prices and adjusted by regional price deflators to the Jakarta price level.

For Peru I use the first (ENAH01, 1997) and third wave (ENAH03, 1999) of the Peruvian Encuesta Nacional de Hogares conducted by the Instituto Nacional de Estadística e Informática. The ENAHO is an ongoing longitudinal living standard measurement survey. It is representative for the three rural and four urban areas of Peru. The 'panel-households' are only a sub-sample of all households interviewed. In total 3,027 households (14,948

⁶For details see Strauss, Beegle, Sikoki *et al.* (2004).

⁷The number of households is higher in the second period, because it includes so called 'split-off' households, i.e. individuals covered by the IFLS1, but who left their initial household and formed their own new household.

individuals) have been followed over the the first three waves. De Vreyer, Mesplé-Soms and Herrera (2002) have shown that there seems to be no significant attrition bias. Attrition could be a problem if the fourth wave (2000) were used, because of a substantial drop out of many panel households. Therefore, I use the third instead of the fourth wave. I use again real household expenditure per capita as the income measure. Expenditure is expressed in 1997 prices and adjusted by regional price deflators to the Lima price level.

3.2 An assessment of pro-poor growth with and without postulating the anonymity axiom

In Indonesia during the first period covered by the IFLS data—1993 to 1997—real GDP per capita increased by almost 5 percent per year. Table 1 shows, as one can expect, that household incomes increased and poverty could be significantly reduced. This very favorable dynamic was abruptly stopped by the economic crisis which started to be felt in the South-East Asia region in April 1997. However the major impact did not hit Indonesia until December 1997/January 1998, just after IFLS2 was conducted. Then, in 1998 , GDP per capita declined almost by 12 percent. The sustained crisis period continued in Indonesia more than a year.⁸ Yet in 2000, when IFLS3 was conducted, the population had—benefiting from the pre-crisis positive dynamic—returned to roughly its pre-crisis living standard, and as Table 1 shows, with some people even a little better off.

[insert Table 1]

A look at the usual (cross-section) GICs (the figures on the left hand side of Figures 1 and 2), which postulate anonymity, show that growth was in both sub-periods positive over the whole income distribution and thus according to the absolute definition ‘pro-poor’. During the period 1993 to 1997 the GIC indicates that growth rates up to the 80th percentile of the income distribution were even higher than the average growth rate and thus growth was also ‘pro-poor’ according to the relative definition. In consequence inequality decreased (see Table 1). This was, except in the first ten percentiles, not the case during the period 1997 to 2000. This can also easily be seen by the fact that during the first period the mean of percentile growth rates was above the growth rate in mean, whereas it was below the growth rate in mean during the second period. Table 2 shows the rates of pro-poor growth, *RPPG*, for both periods and alternative poverty lines. The rates computed under the anonymity axiom (1st and 2nd column) consistently suggest that between 1993 and 1997 growth was highly pro-poor for both poverty lines used and between 1997 and 2000 only ‘moderately’

⁸For details, see e.g. Strauss, Beegle, Dwijanto *et al.* (2002).

pro-poor if the 25 percent poverty line is retained and even negative (or ‘anti-poor’) if the 50 percent poverty line is retained.

[insert Table 2, Figures 1–2]

However, these growth incidence curves completely hide the mobility of individuals and households across the income distribution. They offer only a comparison of marginal distributions and are compatible with various movements of poor and non-poor individuals over time. For instance, one might want to know whether those individuals being poor after the crisis are the same individuals than those being poor before the crisis. Put differently, did post-crisis policies and reforms only help a few poor to escape poverty, or, instead, were these measures very favorable for the poor and helped many of them to improve substantially their living standard, but did in the same time hurt the richer households and pushed some of them below the poverty line? In other words, was growth between 1997 and 2000 really not particularly pro-poor, or was instead mobility very intense, moving many poor out of poverty while pushing others into it? The usual pro-poor growth assessment does not allow to distinguish between both phenomena. From a political point of view, this might of course be crucial. In the first case further policies are needed to attack chronic poverty (e.g. investment in the productivity of the poor). In the second case, poverty seems to be more a transient phenomenon, and policies providing safety nets might be the right response.

To answer these questions, I now turn to the IGICs, i.e. to the growth incidence curves, where growth rates for percentiles containing the same individuals in both years are considered. Looking first at the curve for the period 1993 to 1997 (Figure 1, RHS), one can state that the pattern of the IGIC is even ‘more’ pro-poor than that of the GIC, indicating strong (unconditional) convergence or what is sometimes called ‘regression to the mean’. A look at the other IGICS (Figures 2-5), shows that this ‘regression to the mean’ can be observed more or less for all spells considered. Measurement error might of course be a problem here and be responsible for the observed convergence. However, this problem should, at least partly, be under control, given that the data was trimmed (see Appendix), but it will be addressed in more detail below.

That means with respect to the period 1993 to 1997 the GIC hides the high mobility of individuals over the income distribution and the fact that particularly the initially poor benefited from income growth. This can also be seen when computing the mean growth rate for the 50 percent initially poor (*IRPPG*), which is 16.4 percent instead of the obtained 1.8 percent, when simply the mean growth rate for the lower 50 percent of the income distribution is computed (Table 2). However, in this case both curves show at least qualitatively the same thing: pro-poor growth in the absolute as well as in the relative sense.

The decomposition results presented in Table 3, can describe in more detail the mobility across the income distribution and help to understand how the drawn GIC and IGIC arose. For instance, if the change in inequality (measured by the mean log deviation (*MLD*)) is decomposed into the formation of horizontal inequality, i.e. inequality between initially ‘equals’ and vertical inequality, i.e. inequality between initially ‘unequals’, one can state, that heterogeneity in growth rates for initially equals would have risen the *MLD* by more than 15 points, but this increase in inequality was over-compensated by a significant reduction in inequality between initially unequals (-19 points), i.e. a narrowing of the income distribution. Furthermore, looking at the second decomposition, we state that ‘reranking’ of individuals over the income distribution alone would have risen inequality as measured by the Gini coefficient by more than 18 points. However, this was over-compensated by ‘pro-poor growth’ in the relative sense, i.e. higher growth rates among the poor, indicated by the positive value of the ‘progressivity-component’ *P* of 21 points (which has to be deduced from the reranking component to obtain the total change in inequality). Finally, the decompositions of poverty measures (the 50 percent poverty line is used here) show that the change in the headcount index—the proportion of individuals poor—can be explained by a reduction of 27 points due to ‘leavers’ and an increase of 7 points due to ‘joiners’. Or, if the poverty gap ratio is decomposed, one can state that ‘leavers’ reduced the average distance to the poverty line in relation to the poverty line by almost 9 points, income growth of those staying in poverty reduced it by further 3 points and, in contrast, ‘joiners’ increased it by only 1.6 points. In consequence, the GIC in Figure 1 (LHS) is the result of significant upward mobility and clearly higher growth rates among the poor which reduced also the average gap to the poverty line of those who stayed under the poverty line. This was only slightly be compensated by downward mobility of initially non-poor individuals, and, as a result, vertical inequality decreased.

[insert Tables 3–5]

Making the comparison of the GIC and IGIC for the period 1997-2000, one can state that whereas the GIC is U-shaped, suggesting that for the very poor and the very rich growth was higher than growth in mean, the IGIC has a clear negative slope (again suggesting regression to the mean) and, in contrast to the GIC, growth in mean is significantly below the mean of percentile growth rates. Therefore in this case, whether we postulate or remove the anonymity axiom clearly matters for our conclusion on how the ‘poor’ benefited from growth. This shows also up when computing rates of pro-poor growth (see Table 2). The corresponding rates of pro-poor growth are almost zero for the GIC, but again very high when computed for the IGIC.

Both curves are even more contrasting if they are drawn solely for the urban sample. Whereas postulating anonymity leads to a GIC (Figure 3, LHS) which is clearly anti-poor in the relative sense and only weakly pro-poor in the absolute sense (from the 30th to the 45th percentile), the IGIC (Figure 3, RHS) is clearly pro poor, i.e. growth rates are positive up to the 70th percentile and higher than the growth rate in mean up to the 65th percentile. That means, if we remove the anonymity axiom and consider individual trajectories through time, we get exactly the opposite GIC compared to the case where we do the usual cross-section comparison. Whereas for the GIC the growth rate in mean lies above the mean of percentile growth rates, the opposite is the case for the IGIC. Likewise the rates of pro-poor growth computed for both poverty lines are negative for the GIC, but significantly positive for the IGIC (Table 2).

[insert Figures 2–3]

Again, the decompositions can help to understand what exactly happened to the income distribution and the poor. First one can state that the rise in horizontal inequality was almost exactly offset by the reduction in vertical inequality and likewise that inequality through reranking was completely offset by higher income growth among the poor ($P > 0$). Moreover, the decomposition of the change of the poverty gap ratio shows that incomes from ‘stayers’ stagnated and the reduction in the mean poverty gap due to ‘leavers’ was again exactly compensated by ‘joiners’ having fallen under the poverty line. Hence, whereas the GIC suggests, that the poor did benefit underproportionally from growth, the decomposition shows that almost 30 percent of the poor left poverty, but that this effect was however not reinforced through income growth among ‘stayers’ and even partly off-set through individuals falling under the poverty line. Of course more specific value judgments have to be formulated to decide whether such a growth process should be called ‘pro-poor’ or not.

Now we turn to the Peruvian case. In the nineties Peru had to face substantial institutional reforms and several macro-economic shocks. Among other things the country was adversely affected by the economic crisis in South-East Asia and EL Niño. From 1997 on macro-economic growth slowed down and became even negative in 1998 and 1999.⁹ Table 1 shows that real household income per capita, poverty and inequality remained more or less constant during that period. However, the comparison of the GIC with the IGIC will again show that this ‘cross-sectional’ stability hides interesting dynamics.

Whereas the GIC (Figure 4, LHS) shows positive growth only for the poorest five percentiles and between the 15th and the 25th percentile and negative growth for all others, the IGIC (Figure 4, RHS) indicates positive

⁹For details see e.g. Herrera and Roubaud (2003).

growth rates up to the 75th percentile. As for Indonesia, the slope of the IGIC is clearly negative, indicating higher growth rates for the poor and thus again convergence. Likewise, whereas the mean of percentile growth rates lies below the growth rate in mean for the GIC, it lies not only above the growth rate in mean for the IGIC, but is also positive (about 5 percent). This contrast is even more pronounced if rural areas are considered alone (see Figure 5). On the national level as well as for rural areas, the *RPPGs* are close to zero or even negative, whereas the *IRPPGs* are clearly positive (Table 2).

[insert Figures 4–5]

As for Indonesia in 1997 to 2000, horizontal inequality is almost exactly compensated by a decline of vertical inequality. Likewise, higher inequality due to reranking is one to one offset by the progressivity component, i.e. higher growth rates among the poor. The more or less constancy of the headcount index for the 50 percent poverty line, hides a reduction by 11 points due to ‘leavers’ and an increase by 12 points due to ‘joiners’. Decomposing changes in the poverty gap ratio, yields that ‘leavers’ reduced this poverty measure by almost 3 points, but that this impact was completely over compensated by ‘joiners’ which increased it by more than 3.1 points. Income growth for the chronic poor was very weak and contributed almost nothing to poverty reduction.

3.3 Robustness to measurement error

The above results are all based on a sample of expenditures declared by households (‘income’ in what follows). Apparent outliers have been withdrawn from the sample using the Mahalanobis distance measure (see Appendix). However, it is evident that the remaining declarations are still, more or less, affected by measurement error. When drawing the usual GIC, the measurement error is less a problem given that we only compare marginal income distributions. However, the problem can be more serious, when drawing the IGIC, which is based on a joint income distribution, even if the problem is reduced due to the fact that growth rates are—as for the GIC—computed over percentiles and not over individuals directly. In this sub-section, I will analyze the robustness of the—in more or less all cases stated—negative slope of the IGICs to the existence of measurement error. To do this, I mainly follow the approach suggested by Fields, Cichello, Freije *et al.* (2003).

It is assumed that the income reported by household i in year t is given by the sum of unobserved true income Y_{it}^* and a measurement error component μ_{it} :

$$Y_{it} = Y_{it}^* + \mu_{it}, \quad (10)$$

where μ_{it} may be correlated with true income. Following Fields *et al.* (2003) it is assumed that measurement error in the initial period $t - 1$ is a linear function of true income, plus a white-noise disturbance term, u_{t-1} . If the average true income in the initial period is denoted as \bar{Y}_{t-1}^* and δ_{t-1} represents the correlation between true base year income and measurement error, measurement error in the initial year reported income can be written as:

$$\mu_{it-1} = \delta_{t-1}(Y_{it-1}^* - \bar{Y}_{t-1}^*) + u_{it-1}. \quad (11)$$

Given that measurement error might be correlated over time, a serial correlation coefficient ρ is defined. Measurement error in the final period can then be written as:

$$\mu_{it} = \delta_t(Y_{it}^* - \bar{Y}_t^*) + \rho u_{it-1} + u_{it}. \quad (12)$$

The relationship between households' income in the initial period and their subsequent income change, when income is measured without error, is the coefficient from a regression of true income change on true initial income. This coefficient measures the extent of convergence or divergence in true income and can be expressed as:

$$\beta^* = \frac{\text{Cov}[Y_t^* - Y_{t-1}^*, Y_{t-1}^*]}{\text{Var}[Y_{t-1}^*]}. \quad (13)$$

The OLS estimate from a regression of reported income change on reported base year income is denoted:

$$\beta = \frac{\text{Cov}[Y_t - Y_{t-1}, Y_{t-1}]}{\text{Var}[Y_{t-1}]}. \quad (14)$$

As shown in Fields *et al.* (2003), Equations (10) through (14) now yield:

$$\begin{aligned} \beta = \beta^* \frac{\text{Var}[Y_{t-1}^*]}{\text{Var}[Y_{t-1}]} (1 + \delta_{t-1})(1 + \delta_t) - \frac{\text{Var}[u_{t-1}](1 - \rho)}{\text{Var}[Y_{t-1}]} \\ + \frac{\text{Var}[Y_{t-1}^*]}{\text{Var}[Y_{t-1}]} (1 + \delta_{t-1})(\delta_t - \delta_{t-1}). \end{aligned} \quad (15)$$

To give these three terms an interpretation, two additional assumptions have to be made according to Fields *et al.* (2003). First, a particular household's propensity to misreport income is assumed to decline or remain constant over time, such that $\rho \leq 1$. Second, measurement error is assumed partially correlated with true income, such that δ_{t-1} and δ_t are both > -1 . Both assumptions are consistent with empirical evidence (see Bound, Brown and Mathiowetz 2001).

Under these assumptions, the second term of Equation (15) indicates that the measurement error in initial income contributes to an apparent negative correlation between base-year income and subsequent income change.

This is due to the fact that the measurement error of the initial period enters of course also the computed income change. However, this bias is partly offset if measurement errors are serially correlated. The first term of Equation (15) corresponds to the standard attenuation bias caused by the stochastic independent variable. This attenuation bias is aggravated if measurement error is negatively correlated with true income in each period. As Fields *et al.* (2003) emphasize, this attenuation bias counteracts the effects of the second term by raising the value of β towards zero, whenever the true relationship between initial income and income change is negative, i.e. under convergence. Finally, the third term will be relatively small, unless the correlation coefficient between measurement error and true income changed substantially between periods.

From Equation (15) Fields *et al.* (2003) derive the following expression:

$$\frac{\text{Var}[u_{t-1}]}{\text{Var}[Y_{t-1}^*]} = \frac{\beta^*(1 + \delta_{t-1})(1 + \delta_t) + (1 + \delta_{t-1})(\delta_t - \delta_{t-1}) - \beta(1 + \delta_{t-1})^2}{1 - \rho + \beta}, \quad (16)$$

which gives the variance of stochastic measurement error, relative to the variance of true income, given the observed regression coefficient on reported income β and a particular value of the unknown coefficient on true income, β^* . Setting β^* equal to zero can then give the minimum amount of measurement error required to overturn the negative relationship between initial income and income change.

Table 6 shows the OLS estimates of β , when for each of the five spells analyzed above, $(Y_{it} - Y_{it-1})$ is regressed on Y_{it-1} (without controlling for any other variables, i.e. test of unconditional convergence). Moreover, Table 6 reports the minimum threshold for each spell, which is computed for different combinations of ρ and the δ s (as in Fields *et al.* (2003), it is assumed that δ_{t-1} and δ_t are equal). The chosen parameters ρ and δ correspond to the lower and upper bounds found in various validation studies on earnings declarations summarized in Bound *et al.* (2001). The correlation coefficient between measurement error and true earnings usually seems to lie between -0.05 and -0.4. A reasonable range for serial correlation goes according to these studies from 0.1 to 0.2. These orders of magnitude are derived from declarations on annual earnings and do not necessarily apply to the expenditure data used in this study, but should however, given the wide range of parameters tested, serve as reasonable bounds.

[insert Table 6]

In Indonesia, for divergence to have taken place, the variance of measurement error would need to be at least 75 to 670 percent of the variance of

true incomes, depending on the correlation between measurement error and both true income and past measurement error. In Peru, measurement error with a variance that ranges from 28 to 125 percent of true income could already be responsible for the observed estimates of convergence. Bound *et al.* (2001) report for the ratio of the variance of measurement error to the variance of true income a usual range of 0.1 to 0.3. That means that the observed convergence to the mean and the resulting negatively sloped IGICS can be considered as highly robust against measurement error for the case of Indonesia. However, for the case of Peru, it cannot not be excluded with certainty that measurement error is responsible for the observed convergence.

However, it should be noted that ‘mean-reversion’ ($\beta < 0$) is not sufficient but only necessary to prove convergence, under some circumstances the rate of convergence is even independent of the degree of mean-reversion. Put differently mean-reversion and convergence are, as shown by Lichtenberg (1994), not completely equivalent. But given that in in both countries, Indonesia and Peru, inequality does not substantially rise and for some spells even decrease, the mean reversion should also imply convergence.

4 Conclusion

The assessment of pro-poor growth for Indonesia and Peru with and without postulating the anonymity axiom, has shown that postulating anonymity, that is considering the usual cross-sectional growth incidence curves, completely hides the mobility of individuals across the income distribution and therefore, cannot say much on how exactly a particular growth incidence curve arose. The shape of almost all growth incidence curves constructed using the panel dimension of the data, thus removing anonymity, shows that growth for the initially poor was generally stronger than the usual cross-section growth incidence curve suggested. Put differently, almost each spell considered indicated substantial up-ward mobility suggesting convergence to the mean. In the same time the decomposition of poverty changes showed that (except in Indonesia 1993-1997) a part of the population seems however to be stuck in poverty and to benefit not at all from economic growth (the ‘pro-poor growth component’ is close to zero). In other words, cross-section data cannot be used to track the experience of a particular set of individuals over time, but only to track income groups, whose composition may change. This explains how it can be possible both for ‘the poor’ to fare badly relatively to ‘the rich’ (constant or increasing poverty measures despite growth in mean income) and for income growth to be ‘pro-poor’ (up-ward mobility of initially poor). An extreme case was presented for urban Indonesia for the period 1997 to 2000, where removing anonymity results in an exactly reversed growth incidence curve. It was shown that all these findings are

robust against the influence of measurement error in the case of Indonesia but however not necessarily in the case of Peru.

In any case, the analysis shows that a judgement about the extent of the pro-pooriness of growth based on the usual cross-section growth incidence curve can give a biased impression on how the initially poor benefited or not from growth. In consequence, when postulating anonymity and interpreting growth incidence curves, one should be aware and very explicit on what exactly is measured. To be clear, the objective is not to question the utility of cross-section comparisons of income distributions, but to highlight that they should be complemented by some kind of longitudinal analysis to understand the sources of poverty reduction and inequality, and the inequality and mobility consequences of policies more deeply, i.e. who were the losers and winners of specific reforms. Of course, removing anonymity makes necessary some value judgements on how one weights income mobility vs. changes of cross-section inequality and poverty, i.e. to what extent the initial position of individuals should matter for our judgement about final outcomes of a policy.

Unfortunately, in most cases, especially for developing countries, we do not have panel data at hand and it seems that then we are forced to postulate anonymity when comparing income distributions over time. A solution to this problem can be to rely on micro-simulation methods or some kind of counterfactual analysis. Micro-simulation became recently quite popular in development economics due to the *PRSP Initiative*¹⁰ and the resulting need to evaluate income distributional consequences of macro-economic shocks and policy reforms. The principle of this methodology is to estimate, mostly econometrically, some kind of income model $y_i = Y(x_i, \epsilon_i; \beta; \sigma; \lambda) / P(c_i; p)$, where real household income per capita, y_i , of household i is assumed to depend on six sets of arguments: its observable socio-demographic characteristics, or those of its members (x_i), unobservable characteristics (ϵ_i), a vector of remuneration rates of the observed (β) and unobserved earnings determinants (σ), and a set of parameters defining the participation and occupational choice behavior of its members (λ_t). P stands for a household specific price deflator being a function of the household's budget shares c_i and a vector of commodity prices p . This model is then used to simulate counterfactual incomes \tilde{y}_i by varying in accord with historical observations or some specific policy reform either one or several of the variables and/or one or several of the parameters. Finally, this allows to analyze the joint distribution $F(y_i, \tilde{y}_i)$. For a detailed description of this kind of methodology see e.g. Bourguignon and Ferreira (2003) and Cogneau, Grimm and Robilliard (2003).

The only limited use of comparing marginal distributions when evaluating specific policies was recently also highlighted by Cunha, Heckman and

¹⁰'PRSP' stands for *Poverty Reduction Strategy Paper*.

Navarro (2004).¹¹ They suggested a similar methodology to that described above allowing to recover the joint distribution $F(y_0, y_1)$ if for each individual depending on her ‘treatment’, $s = 0$ or $s = 1$, only y_0 or y_1 is observed. The principle of this methodology is to analyze econometrically choice data using a factor structure model in order to construct for each individual the counterfactual income \tilde{y}_0 if $s = 1$ and \tilde{y}_1 if $s = 0$. Once the whole counterfactual income distribution is recovered a social mobility analysis can be undertaken.

Obviously, these methodologies are more time intensive than simple income distribution analyses. However, they have the advantage that they can also solve the usual problem inherent in ‘before-after-comparisons’ by isolating what distributional change is due to a specific shock or policy and what is due to other changes. Thus, these methodologies eventually allow to go also beyond to what can be done with panel data analysis.

Appendix: Elimination of outliers

To eliminate the influence of outliers the data were trimmed. I use a similar method to that used by Jenkins and Van Kerm (2003). For each pair of years analyzed, I discarded an observation if the *Mahalanobis* distance between its two log per capita income values (y_{it-1}, y_{it}) exceeded a critical value equal to the mean plus two times the standard deviation of the distribution of the *Mahalanobis* distances in the two-year sample. The vector of the household’s *Mahalanobis* distances for each pair of years can be computed by the following equation:

$$MD = \left[\text{diag}[(y - m)S^{-1}(y - m)'] \right]^{1/2} \quad (17)$$

where y is a $N \times 2$ matrix containing the two vectors of incomes of both years, m is a 1×2 matrix containing the mean incomes of both years, and S is the 2×2 variance-covariance matrix of the two income vectors. Income is in logarithmic terms and expressed on a per capita basis.

The advantage of using this concept is that the *Mahalanobis* distance identifies not only outlier incomes in each year but also outlier changes in income between years. Applying this concept, between 4 and 5 percent of the observations in each two-year sample were—besides some few observations, were not any information at all on expenditures was available—excluded.

References

- Atkinson A.B. (1970), On the Measurement of Inequality. *Journal of Economic Theory*, 2: 244-63.

¹¹See also Carneiro, Hansen and Heckman (2001).

- Auerbach A.J. and K.A. Hassett (2002), A New Measure of Horizontal Equity. *American Economic Review*, 92 (4): 1116-1125.
- Bound J., C. Brown, and N. Mathiowetz (2001), Measurement error in survey data. In J.J. Heckman and E. Learner (eds.), *Handbook of Econometrics* Vol. 5 (pp. 3705-3843), Amsterdam: Elsevier Science.
- Bourguignon F. and F.H.G. Ferreira (2003), Ex Ante Evaluation of Policy Reforms Using Behavioral Models. In F. Bourguignon and L. Perreira Da-Silva (eds.), *The Impact of Economic Policies on Poverty and Income Distribution* (pp. 123-141), Washington D.C.: World Bank.
- Carneiro P., K.T. Hansen and J.J. Heckman (2001), Removing the Veil of Ignorance in assessing the distributional Impacts of Social Policies. *Swedish Economic Policy Review*, 8: 273-301.
- Cleveland W.S. (1979), Robust Locally Weighted Regression and Smoothing Scatter Plots. *Journal of the American Statistical Association*, 74: 829-836.
- Cogneau D., M. Grimm and A.-S. Robilliard (2003), Evaluating poverty reduction policies — The contribution of micro-simulation techniques. In J.-P. Cling, M. Razafindrakato and F. Roubaud (eds.), *New International Poverty Reduction Strategies* (pp. 340-370), London: Routledge Books.
- Cowell F.A. (2000), Measurement of Income Inequality. In A.B. Atkinson and F. Bourguignon (eds.), *Handbook of Income Distribution* (pp. 87-166), Amsterdam: North-Holland.
- Cunha F., J. Heckman and S. Navarro (2004), Counterfactual Analysis of Inequality and Social Mobility. Mimeo, University of Chicago, Department of Economics, Chicago.
- De Vreyer P., S. Mesplé-Somps and J. Herrera (2002), Consumption growth and spatial poverty traps: an analysis of the effects of social services and community infrastructures on living standards in rural Peru. DIAL Working Paper DT/2002/17, DIAL, Paris.
- Duclos J.-Y. and P.J. Lambert (2000), A Normative Approach to Measuring Classical Horizontal Inequity. *Canadian Journal of Economics*, 33: 87-113.
- Duclos J.-Y. and Q. Wodon (2004), What is “Pro-Poor”?, Mimeo, World Bank, Washington D.C. and CIRPÉE, Université Laval, Québec.
- Foster J., J. Greer and E. Thorbecke (1984), A Class of Decomposable Poverty Measures. *Econometrica*, 52: 761-776.

- Hicks J.R. (1939), The Foundations of Welfare Economics. *Economic Journal*, 69: 696-712.
- Fields G.S. and E.A. Ok (1996), The Meaning and Measurement of Income Mobility. *Journal of Economic Theory*, 71 (2): 349-377.
- Fields G.S., P.L. Cichello, S. Freije, M. Menéndez and D. Newhouse (2003), For Richer or for Poorer? Evidence from Indonesia, South-Africa, Spain, and Venezuela. *Journal of Economic Inequality* 1: 67-99.
- Herrera J. and F. Roubaud (2003), Dynamique de la pauvreté urbaine au Pérou et à Madagascar 1979-1999: Une analyse sur données de panel. DIAL Working Paper DT/2003/03, DIAL, Paris.
- Hulme D. and A. Shepard (2003), Conceptualizing Chronic Poverty. *World Development*, 31 (3): 403-424.
- Jenkins S.P. and P.J. Lambert (1999), Horizontal Inequality Measurement: A Basic Reassessment. In J. Silber (ed.), *Handbook of Income Inequality Measurement*, Boston: Kluwer Academic Publishers.
- Jenkins S.P. and P. Van Kerm (2003), Trends in Income Inequality, Pro-Poor Income Growth and Income Mobility. IRISS Working Paper No. 2003-11, CEPS/Instead, Luxembourg.
- Kakwani N. and E. Pernia (2000), What is Pro-Poor Growth? *Asian Development Review*, 18: 1-16.
- Kakwani N., S. Khandker and H.H. Son (2003), Poverty Equivalent Growth Rate: With Applications to Korea and Thailand. Mimeo, Economic Commission for Africa.
- Kaldor N. (1939), Welfare Propositions in Economics and Interpersonal Comparisons of Utility. *Economic Journal*, 69: 549-552.
- Kaplow L. (1989), Horizontal Equity: Measures in Search of a Principle. *National Tax Journal*, 42: 139-154.
- Klasen S. (2004), In Search of the Holy Grail. How to Achieve Pro Poor Growth. In N. Stern. and B. Tungodden (eds.), *Towards Pro Poor Policies. Proceedings of the Annual Bank Conference on Development Economics (ABCDE) Europe Conference.*, Washington D.C.: World Bank.
- Lichtenberg F.R. (1994), Testing the Convergence Hypothesis. *Review of Economics and Statistics*, 76 (3): 576-579.
- McCulloch N. and B. Baulch (2000), Tracking pro-poor growth. ID21 insights No.31. University of Sussex, Institute of Development Studies.

- Plotnick R. (1982), The Concept and Measurement of Horizontal Inequity. *Journal of Public Economics*, 17: 373-391.
- Rawls J. (1971), *A theory of justice*. Cambridge, MA: Cambridge University Press.
- Ravallion M. (2004a). Pro-Poor Growth: A Primer. World Bank Policy Research Working Paper No. 3242, World Bank, Washington D.C.
- Ravallion M (2004b), Competing Concepts of Inequality in the Globalization Debate. World Bank Policy Research Working Paper No. 3243, World Bank, Washington D.C.
- Ravallion M. and S. Chen (2003). Measuring Pro-Poor Growth. *Economics Letters*, 78: 93-99.
- Ravallion M. and M. Lokshin (2004), Gainers and Losers from Agricultural Trade Reform in Morocco. World Bank Policy Research Working Paper No. 3368, World Bank, Washington D.C.
- Strauss J., K. Beegle, A. Dwiyanto, Y. Herawati, D. Pattinasarany, E. Satriawan, B. Sikoki, B. Sukamdi and F. Witoelar (2002), Indonesian living standards three years after the crisis: evidence from the Indonesia Family Life Survey, Executive Summary, RAND Corporation.
- Strauss J., K. Beegle, B. Sikoki, A. Dwiyanto, Y. Herwati and F. Witoelar (2004), The Third Wave of the Indonesia Family Life Survey (IFLS3): Overview and Field Report. WR-144/1-NIA/NCHID, RAND Corporation.

Tables and Figures

Table 1
Growth, poverty and inequality

	Initial	Final
<i>Indonesia, 1993–1997, national</i>		
Growth in mean p.y.		0.079
FGT0 (25% pline.)	0.250	0.100
FGT1 (25% pline.)	0.072	0.023
FGT0 (50% pline.)	0.500	0.297
FGT1 (50% pline.)	0.185	0.086
Gini-Coeff.	0.400	0.376
<i>Indonesia, 1997–2000, national</i>		
Growth in mean p.y.		0.019
FGT0 (25% pline.)	0.250	0.228
FGT1 (25% pline.)	0.069	0.058
FGT0 (50% pline.)	0.500	0.473
FGT1 (50% pline.)	0.173	0.158
Gini-Coeff.	0.363	0.367
<i>Indonesia, 1997–2000, urban only</i>		
Growth in mean p.y.		0.011
FGT0 (25% pline.)	0.250	0.252
FGT1 (25% pline.)	0.067	0.070
FGT0 (50% pline.)	0.500	0.502
FGT1 (50% pline.)	0.177	0.179
Gini-Coeff.	0.354	0.372
<i>Peru, 1997–1999, national</i>		
Growth in mean p.y.		-0.008
FGT0 (25% pline.)	0.250	0.247
FGT1 (25% pline.)	0.071	0.069
FGT0 (50% pline.)	0.500	0.514
FGT1 (50% pline.)	0.191	0.194
Gini-Coeff.	0.367	0.366
<i>Peru, 1997–1999, rural only</i>		
Growth in mean p.y.		-0.009
FGT0 (25% pline.)	0.250	0.256
FGT1 (25% pline.)	0.054	0.056
FGT0 (50% pline.)	0.500	0.507
FGT1 (50% pline.)	0.161	0.167
Gini-Coeff.	0.325	0.327

Notes: FGT0 refers to the headcount index, i.e. the percentage of poor individuals. FGT1 refers to the poverty gap ratio, i.e. the average distance between income and the poverty line (where for non-poor households this distance is set to zero) as a fraction of the poverty line. Two poverty lines are used: the first considers the first 25 percent (25% pline.) and the other considers the first 50 percent (50% pline.) at the bottom of the income distribution in each base year as poor.

Source: IFLS1, IFLS2, IFLS3, ENAHO1, ENAHO3; computations by the author.

Table 2
Rates of pro-poor growth
with and without postulating anonymity

	25% pline.		50% pline.	
	anonymity (<i>RPPG</i>)	no anonym. (<i>IRPPG</i>)	anonymity (<i>RPPG</i>)	no anonym. (<i>IRPPG</i>)
Indo., 1993–1997, national	0.023	0.229	0.018	0.164
Indo., 1997–2000, national	0.007	0.200	-0.002	0.131
Indo., 1997–2000, urban	-0.007	0.167	-0.003	0.126
Peru, 1997–1999, national	0.007	0.200	-0.002	0.131
Peru, 1997–1999, rural	-0.002	0.224	-0.006	0.138

Notes: The rates of pro-poor growth are computed for two alternative poverty lines: the first considers the first 25 percent (25% pline.) and the other considers the first 50 percent (50% pline.) at the bottom of the income distribution in each base year as poor.

Source: IFLS1, IFLS2, IFLS3, ENAHO1, ENAHO3; computations by the author.

Table 3
Decomposition of changes in inequality
in changes in horizontal and vertical inequality

	Mean log deviation			Decomposition	
	Initial	Final	Change	Horizontal	Vertical
Indo., 1993–1997, national	0.268	0.235	-0.033	0.157	-0.190
Indo., 1997–2000, national	0.219	0.221	0.003	0.145	-0.143
Indo., 1997–2000, urban	0.207	0.230	0.023	0.145	-0.122
Peru, 1997–1999, national	0.226	0.222	-0.003	0.100	-0.103
Peru, 1997–1999, rural	0.170	0.172	0.001	0.118	-0.116

Source: IFLS1, IFLS2, IFLS3, ENAHO1, ENAHO3; computations by the author.

Table 4
Decomposition of changes in inequality in changes
due to reranking and changes due to pro-poor growth

	Gini-coefficient			Decomposition	
	Initial	Final	Change	Reran.	Pro-p. growth
Indo., 1993–1997, national	0.400	0.376	-0.024	0.185	0.210
Indo., 1997–2000, national	0.363	0.367	0.005	0.178	0.174
Indo., 1997–2000, urban	0.354	0.372	0.019	0.182	0.163
Peru, 1997–1999, national	0.367	0.366	-0.001	0.100	0.101
Peru, 1997–1999, rural	0.325	0.327	0.001	0.124	0.123

Source: IFLS1, IFLS2, IFLS3, ENAHO1, ENAHO3; computations by the author.

Table 5
Decomposition of changes in poverty
in mobility and pro-poor growth components

	FGT0 (50% pline.)			Decomposition		
	Initial	Final	Change	Out-mob.	Pro-p. growth	In-mob.
Indo., 1993–1997, national	0.500	0.297	-0.202	-0.270	0	0.067
Indo., 1997–2000, national	0.500	0.473	-0.028	-0.158	0	0.131
Indo., 1997–2000, urban	0.500	0.502	0.003	-0.138	0	0.140
Peru, 1997–1999, national	0.500	0.514	0.014	-0.106	0	0.120
Peru, 1997–1999, rural	0.500	0.507	0.008	-0.131	0	0.138
	FGT1 (50% pline.)			Decomposition		
	Initial	Final	Change	Out-mob.	Pro-p. growth	In-mob.
Indo., 1993–1997, national	0.185	0.086	-0.099	-0.087	-0.029	0.016
Indo., 1997–2000, national	0.173	0.158	-0.016	-0.045	-0.059	0.036
Indo., 1997–2000, urban	0.177	0.178	0.001	-0.036	0.001	0.037
Peru, 1997–1999, national	0.191	0.194	0.003	-0.026	-0.001	0.031
Peru, 1997–1999, rural	0.161	0.167	0.005	-0.036	0.004	0.038
	Watts (50% pline.)			Decomposition		
	Initial	Final	Change	Out-mob.	Pro-p. growth	In-mob.
Indo., 1993–1997, national	0.266	0.115	-0.151	-0.120	-0.053	0.021
Indo., 1997–2000, national	0.243	0.216	-0.027	-0.061	-0.012	0.047
Indo., 1997–2000, urban	0.247	0.251	0.004	-0.047	0.003	0.048
Peru, 1997–1999, national	0.273	0.275	0.002	-0.034	0.003	0.040
Peru, 1997–1999, rural	0.215	0.221	0.006	-0.047	0.004	0.049

Notes: The used poverty line (50% line) refers to the poverty line which considers the first 50 percent at the bottom of the income distribution in each base year as poor.

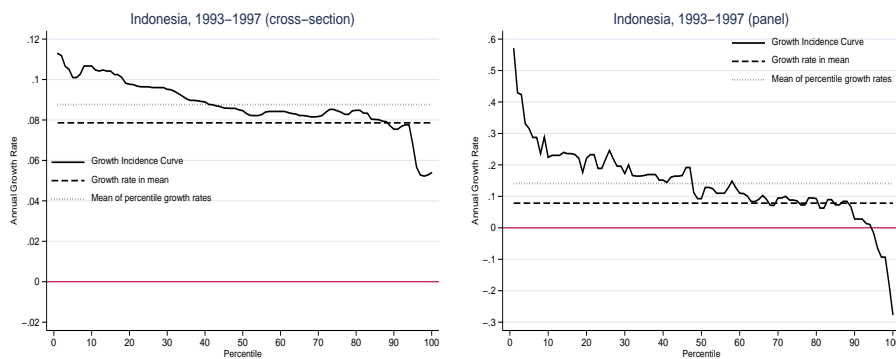
Source: IFLS1, IFLS2, IFLS3, ENAHO1, ENAHO3; computations by the author.

Table 6
 Ratio of measurement error to true income variance
 implying zero correlation between true initial income and true income change

		Indonesia			Peru	
δ	ρ	1993-1997, nat. $\beta = -0.684$	1997-2000, nat. $\beta = -0.686$	1997-2000, urb. $\beta = -0.696$	1997-1999, nat. $\beta = -0.442$	1997-1999, rur. $\beta = -0.441$
0	0	2.165	2.185	2.289	0.792	0.789
0	0.1	3.167	3.206	3.412	0.965	0.961
0	0.2	5.897	6.018	6.692	1.235	1.228
-0.1	0	1.753	1.770	1.854	0.642	0.639
-0.1	0.1	2.565	2.597	2.764	0.782	0.778
-0.1	0.2	4.776	4.874	5.421	1.000	0.995
-0.2	0	1.385	1.398	1.465	0.507	0.505
-0.2	0.1	2.027	2.052	2.184	0.618	0.615
-0.2	0.2	3.774	3.851	4.283	0.790	0.786
-0.4	0	0.779	0.786	0.824	0.285	0.284
-0.4	0.1	1.140	1.154	1.228	0.347	0.346
-0.4	0.2	2.123	2.166	2.409	0.444	0.442

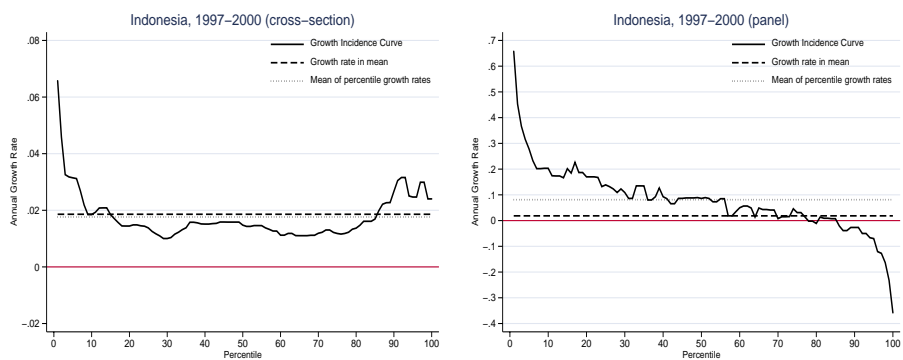
Source: Computations by the author.

Figure 1
 Growth incidence curves: Indonesia, 1993–1997, national
 LHS: anonymity (GIC), RHS: no anonymity (IGIC)



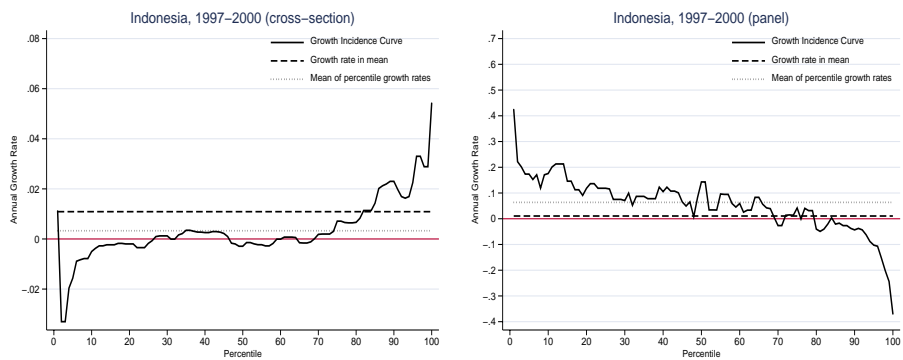
Notes: Curves are smoothed by three period running medians.
 Source: IFLS1, IFLS2, IFLS3; computations by the author.

Figure 2
 Growth incidence curves: Indonesia, 1997–2000, national
 LHS: anonymity (GIC), RHS: no anonymity (IGIC)



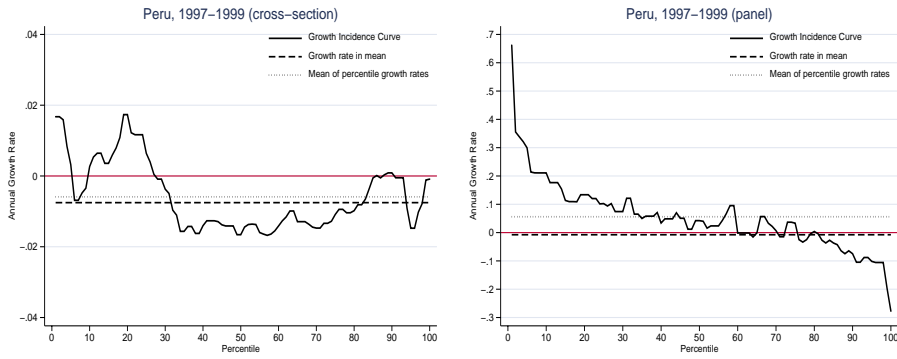
Notes: Curves are smoothed by three period running medians.
 Source: IFLS1, IFLS2, IFLS3; computations by the author.

Figure 3
 Growth incidence curves: Indonesia, 1997–2000, urban
 LHS: anonymity (GIC), RHS: no anonymity (IGIC)



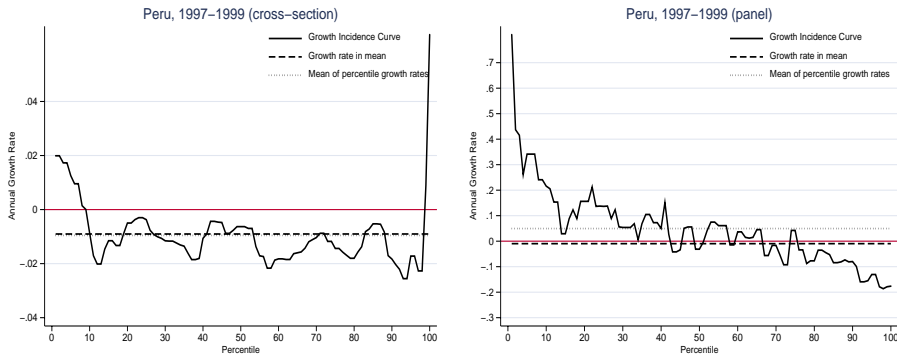
Notes: Curves are smoothed by three period running medians.
 Source: IFLS1, IFLS2, IFLS3; computations by the author.

Figure 4
 Growth incidence curves: Peru, 1997–1999, national
 LHS: anonymity (GIC), RHS: no anonymity (IGIC)



Notes: Curves are smoothed by three period running medians.
 Source: ENAHO1, ENAHO3; computations by the author.

Figure 5
 Growth incidence curves: Peru, 1997–1999, rural
 LHS: anonymity (GIC), RHS: no anonymity (IGIC)



Notes: Curves are smoothed by three period running medians.
 Source: ENAHO1, ENAHO3; computations by the author.