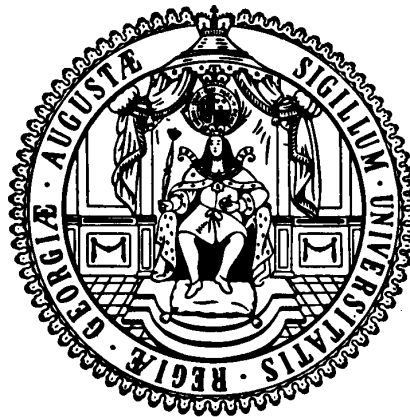


**Ibero-Amerika Institut für Wirtschaftsforschung
Instituto Ibero-Americano de Investigaciones Económicas
Ibero-America Institute for Economic Research
(IAI)**

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



Diskussionsbeiträge · Documentos de Trabajo · Discussion Papers

Nr. 160

Robust Multiperiod Poverty Comparisons

**Johannes Gräßl
Michael Grimm**

July 2007

Robust Multiperiod Poverty Comparisons

Johannes Gräb* and Michael Grimm†

Third draft: July 4, 2007

Abstract

We propose a new methodology for comparing poverty over multiple periods across time and space that does not arbitrarily aggregate income over various years or rely on arbitrarily specified poverty lines or poverty indices. Following Duclos *et al.* (2006a), we use the multivariate stochastic dominance methodology to create dominance surfaces for different time spans. We elaborate the method first for the bidimensional case, using as dimensions income observed over two periods: one at the beginning and one at the end of a time span. Subsequently, we extend it to the case where incomes are observed over n -periods. We illustrate our approach by performing poverty comparisons using data for Indonesia and Peru.

Key words: Multiperiod Poverty, Poverty Dominance, Poverty Dynamics, Chronic Poverty.

*University of Göttingen, Center for Statistics and Department of Economics, Platz der Göttinger Sieben 3, 37073 Göttingen, jgraeb@uni-goettingen.de.

† Institute of Social Studies, Kortenaerkade 12, P.O.Box 29776, NL - 2502 LT The Hague, University of Göttingen, DIW Berlin and DIAL Paris, email: grimm@iss.nl.

We are particularly grateful to Ravi Kanbur for inspiring discussions and to Stefan Dercon and James Foster for very useful comments and suggestions on a first draft of this paper. We also thank Jean-Yves Duclos and Abdelkrim Araar for having shared their Stata program for calculating stochastic dominance over multiple dimensions. Finally, we thank Javier Herrera and Sandrine Mesplé-Somps for having made available their Peruvian household data. Financial support for this research from the Chronic Poverty Research Center at the University of Manchester is greatly acknowledged.

1 Introduction

Today it is well established that poverty is a dynamic phenomenon. But if poverty does fluctuate and evolve over time, this raises the question of how best to measure it over multiple periods. Cross-sectional poverty measures can provide abundant information on the extent of poverty at a given point, but almost none on the rate at which people escape from or fall into poverty over time.

Recognizing this, authors such as Kanbur and Grootaert (1995) have suggested focusing on households' changes in poverty status. Others have developed concepts to aggregate incomes over multiple periods (i.e., trajectories of income over time) using an evaluation function that explicitly captures, for example, the risk aversion of households (see e.g., Cruces, 2005). While such an approach has the advantage of accounting for the negative effects of income variability on the well-being of households, it requires arbitrary assumptions about how exactly 'risk-adjusted mean income' is best computed.

Likewise, considering the various methodologies proposed for measuring and conceptualizing chronic and transient poverty, one can safely state that the results and consequently the policy implications depend heavily on how the two forms of poverty are measured: how incomes are aggregated over time, how the poverty line is set, and what poverty index is chosen (see, e.g., Hulme and McKay, 2005; Jalan and Ravallion 1998; Duclos, Araar and Giles, 2006). Many of these measures are based on some calculation of average income over time and thus abstract from the exact pattern of the income trajectory. In other words, three years of high income followed by three years of low income are treated as six years, with each year of high income followed by a year of low income. In addition, such measures often use one specific poverty line and one specific poverty function.

To circumvent these problems, we suggest a new approach for multi-period poverty measurement based on stochastic dominance tests. This enables us to establish poverty orderings that are valid for a wide range of aggregation rules of incomes observed over time, a wide range of poverty indices, and a wide range of poverty lines. Our approach is inspired by the literature on multi-dimensional poverty orderings (Duclos, Sahn and Younger, 2006a), where dimensions refer to various indicators of individual well-being such as income, education and health.¹ Our dimensions are incomes observed at different points in time. Defining dimensions in this way raises some further challenges, which we discuss below. We develop our approach first for the case where incomes are observed over two periods and then extend it to the case where incomes are observed over n -periods. We

¹See also Duclos, Sahn and Younger (2006b) and the seminal papers by Bourguignon and Chakravarty (2002, 2003).

illustrate our approach using longitudinal data for Indonesia and Peru.

The remainder of our paper is organized as follows. In Section 2 we present our methodology. In Section 3 we provide an empirical illustration. In Section 4 we discuss the results and conclude.

2 Methodology

2.1 Stochastic dominance in a one-period welfare measure

Tests of stochastic dominance are today widely used to establish poverty orderings $\mathbf{P}(\mathbf{Z})$ that are robust for a broad class of poverty measures, $P(F; z)$, and a large range of poverty lines, $z \in Z$. If the welfare measure is denoted y_t and $F(y_t)$ its cumulative distribution function, then stochastic dominance, states that

$$F(y_1)\mathbf{P}(\mathbf{Z})F(y_2) \text{ if and only if } P(F(y_1); z) \leq P(F(y_2); z) \forall z \in Z$$

$$\text{and } P(F(y_1); z) < P(F(y_2); z) \text{ for some } z \in Z, \quad (1)$$

where $F(y_1)\mathbf{P}(\mathbf{Z})F(y_2)$ means that $F(y_1)$ has unambiguously less poverty than $F(y_2)$ with respect to the poverty index P and the range Z .

This result holds for all poverty measures within the Foster-Greer-Thorbecke family, P_α (Foster, Greer and Thorbecke, 1984) with $\alpha \geq 0$ (Foster and Shorrocks, 1988a, b).²

The concept of poverty dominance is useful because it circumvents the problem of choosing one particular poverty measure and one specific poverty line. In the following, we extend the concept, first to two-period welfare measures and then to n -period measures.

2.2 Stochastic dominance in a two-period welfare measure

Assuming that the dynamics of poverty are important, we do not use a one-period welfare measure y_t to define poverty but rather a two-period welfare measure (y_1, y_2) . Furthermore, we impose the condition that well-being is differentiable with respect to the welfare measure in $t = 1$ and $t = 2$ and that income in both periods contributes positively to individual well-being. However, we impose nothing regarding the precise value of the contribution of each year to individual well-being.

Obviously, as is the case for period-by-period poverty comparisons, it is desirable that poverty comparisons over multiple time spans, T_j , like for

²The Foster-Greer-Thorbecke poverty measure has the formula $P_\alpha = 1/N \sum_{i=1}^z (\frac{z-y_i}{z})^\alpha$, where N is the total number of individuals $i = 1, \dots, N$. The parameter $\alpha > 0$ is a poverty aversion parameter: $\alpha = 0$ yields the poverty headcount index, $\alpha = 1$ the poverty gap index, and $\alpha = 2$ poverty severity index (Foster, Greer and Thorbecke, 1984).

example $T_A = [t = 1a; t = 2a]$ vs. $T_B = [t = 1b; t = 2b]$, where t refers to periods (e.g. years), are robust to a large set of poverty lines $z \in Z$. This can be ensured by simply transferring the concept of stochastic dominance for univariate welfare distributions to the case of bivariate welfare distributions.

Furthermore, poverty orderings in the bivariate case should be robust to a broad range of procedures for aggregating the observed welfare measures in the two periods constituting a time span. Thus, the weight attributed to each single period should not matter, e.g., whether we discount income in the second period of the time span to the present value in the first period, or whether we instead weigh income higher in the second period relative to the first. Discounting seems particularly sensible when the time spans under comparison differ in length. Assigning different weights to different periods can also be useful when observing income over different seasons within a time span. In a developing country, prices usually have high seasonal fluctuations. The differences in prices (beyond the usual annual inflation) as well as the disutility stemming from price-induced substitution effects could be incorporated by weighting the first and second periods differently.

However, there is one special situation in which robustness to the aggregation procedure cannot be guaranteed. This arises when the time spans under consideration overlap, i.e., when the second period of the first time span simultaneously represents the first period of the second time span. In this case, the same weight has to be assigned to each. This problem obviously cannot occur with comparisons over space, which are always robust to the aggregation procedure.

Considering two time spans T_A and T_B of equal length and without overlap, we say that

$$F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b})$$

if and only if $P(F(y_{1a}, y_{2a}); z_1, z_2) \leq P(F(y_{1b}, y_{2b}); z_1, z_2) \forall z_1, z_2 \in Z$

and $P(F(y_{1a}, y_{2a}); z_1, z_2) < P(F(y_{1b}, y_{2b}); z_1, z_2)$ for some $z_1, z_2 \in Z$, (2)

where $F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b})$ means that multiperiod poverty is lower over time span T_A than over time span T_B with respect to the poverty index P , the range Z and any aggregation procedure of incomes observed in the two periods constituting a time span.

In this most general case, our concept makes it possible to choose $z_1 \neq z_2$ s.t. $z_{1a} = z_{1b}$ and $z_{2a} = z_{2b}$, i.e., to give a different weight to the first and second period within each time span. Again, such weights could reflect risk aversion or a preference for the present as mentioned above. In this case the test domain for dominance represents a rectangle, where $y_1 < z_1 \wedge y_2 < z_2$. This is illustrated in Figure ???. Note that choosing $z_1 \neq z_2$ is one way of giving different weights to period 1 and 2. Alternatively, we could define explicitly weights for periods 1 and 2, for example, w_1 and w_2 , or apply a discount rate r to period 2.

[insert Figure 1 about here]

In our methodology, and in contrast to ‘one-period-stochastic-dominance’ $F(y_1, y_2)$ refers now to a bivariate distribution. Hence, the test of stochastic dominance does not imply comparing two curves, as with one-period welfare measures, but two surfaces, where each surface is characterized by its two dimensions—the welfare measure in the first and second period—and the cumulative density at each point of that surface.

In the special case where the time spans have an overlap, i.e., $y_{2a} = y_{1b}$, and, hence, an equal weight has to be given to each period (i.e., we use the same poverty line for period 1 and period 2), the dominance criteria simplifies to

$$\begin{aligned}
 & F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b}) \\
 & \text{if and only if } P(F(y_{1a}, y_{2a}); z) \leq P(F(y_{1b}, y_{2b}); z) \quad \forall z \in Z \\
 & \text{and } P(F(y_{1a}, y_{2a}); z) < P(F(y_{1b}, y_{2b}); z) \text{ for some } z \in Z, \quad (3)
 \end{aligned}$$

where $F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b})$ means that $F(y_{1a}, y_{2a})$ has unambiguously less poverty than $F(y_{1b}, y_{2b})$ with respect to the poverty index P and the range Z , i.e., multiperiod poverty is less over time span T_A than over time span T_B . Note that we now only test dominance between the two surfaces along an expansion path of z , where $y_1 < z \wedge y_2 < z$.

Our concept is similar to the concept of multi-dimensional poverty comparisons suggested by Duclos *et al.* (2006a), where the dimensions refer not to different periods but to different dimensions of human well-being, such as income, education and health. However, Duclos *et al.* (2006a) make stronger assumptions about the aggregation procedure for the various dimensions by assuming that the poverty measure is entered in a multiplicative way, meaning that the marginal poverty benefit of an increase in one dimension decreases with the value of the other dimension. Put differently, the more someone has in one dimension, the less overall poverty is deemed to decrease if well-being in the other dimension increases.

2.3 Stochastic dominance in a n -period welfare measure

Extending our methodology to the n -period case is straightforward. We say

$$\begin{aligned}
 & F(y_{1a}, y_{2a}, \dots, y_{na})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b}, \dots, y_{nb}) \\
 & \text{if and only if } P(F(y_{1a}, y_{2a}, \dots, y_{na}); z_1, z_2, \dots, z_n) \leq \\
 & P(F(y_{1b}, y_{2b}, \dots, y_{nb}); z_1, z_2, \dots, z_n) \quad \forall z_1, z_2, \dots, z_n \in Z \\
 & \text{and } P(F(y_{1a}, y_{2a}, \dots, y_{na}); z_1, z_2, \dots, z_n) < P(F(y_{1b}, y_{2b}, \dots, y_{nb}); z_1, z_2, \dots, z_n) \\
 & \text{for some } z_1, z_2, \dots, z_n \in Z, \quad (4)
 \end{aligned}$$

where $F(y_{1a}, y_{2a}, \dots, y_{na})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b}, \dots, y_{nb})$ means that multiperiod poverty is less over time span T_A than over time span T_B with respect to the poverty index P , the range Z and any aggregation procedure of incomes observed in the n periods constituting a time span.

Thus, in this most general case, our concept allows us again to choose $z_1 \neq z_2 \dots z_{n-1} \neq z_n$ s.t. $z_{1a} = z_{1b}, z_{2a} = z_{2b}, \dots, z_{na} = z_{nb}$, i.e., to give a different weight to the n periods within each time span. $F(y_1, y_2, \dots, y_n)$ refers now to a n -variate distribution and, hence, the test of stochastic dominance now implies comparing two surfaces, where each surface is characterized by its n dimensions—the welfare measure observed over the n periods—and the cumulative density at each point of that surface.

An additional issue that arises in the n period case is how exactly the two time spans are compared. Theoretically, one can compare time spans built using all possible permutations of periods as long as each time span has the same number of periods and as long as the beginning and the end of the first time span each precede the beginning and the end of the second time span respectively. One can then even test for dominance over all these comparisons. Below we illustrate such a case using time spans of a maximal length of four years.

2.4 Estimation and inference

To establish dominance empirically, it is sufficient—as was shown by Duclos *et al.* (2006a)—to calculate the differences of $\hat{F}(y_{1a}, y_{2a}, \dots, y_{na})$ and $\hat{F}(y_{1b}, y_{2b}, \dots, y_{nb})$ on a sufficiently narrow grid of test points and to test the statistical significance of these differences based on student t -tests.

2.5 Bounds to multidimensional dominance

When applying the methodology presented above, one needs to define a maximum poverty set $\lambda^*(z_1, z_2, \dots, z_n \in Z)$. Obviously, defining that frontier is always arbitrary. We again follow Duclos *et al.* (2006a) and estimate that frontier directly from our sample as the maximum λ^+ for which multiperiod poverty dominance holds. Then we can locate within λ^+ all of the possible poverty frontiers for which there is necessarily more poverty in time span A than in time span B . We then can judge on a case-by-case basis whether these critical sets and frontiers are wide enough to justify the conclusion of poverty dominance.

3 Empirical illustration

3.1 Data

To illustrate the methodology presented above, we use longitudinal data for Indonesia and Peru.

For Indonesia, we 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 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 have moved are tracked to their new location and, where possible, interviewed there (for details see Strauss, Beegle, Sikoki *et al.* (2004)). Using the three waves, we built two panels one from 1993 to 1997 and one from 1997 to 2000, each comprising roughly 32,000 individuals living in 7,000 households. We use real household expenditure per capita as the welfare measure, but refer to it as income in the following. Expenditure is expressed in 1993 prices and adjusted by regional price deflators to the Jakarta price level.

For Peru we use six waves (1997-2002) of the yearly Peruvian Encuesta Nacional de Hogares conducted by the Instituto Nacional de Estadística e Informática. The ENAHO is representative for the three rural and four urban areas of Peru. The 'panel-households' are only a sub-sample of all households interviewed. Each year, some households drop out of the panel and others are added (rotating panel). We construct several year-to-year panels, each containing, with a few exceptions, more than 5,000 individuals living in more than 1000 households. We use again real household expenditure per capita as the income measure. Expenditure is expressed in 2002 prices and adjusted by regional price deflators to the Lima price level.

Given that we use both data sets purely for illustrative purposes, we do not discuss here the potential bias in our analysis through attrition of households or through any other change in the household survey design over time.

Data on Purchasing Power Parity (PPP) over GDP has been taken from the Penn World Table 6.1 (see Heston *et al.* (2002)). In PWT 6.1 the base year is 1996. For our illustration, we compute the PPPs for alternative years by applying the relative rates of inflation to 1996 between the base country, the United States, and Peru and Indonesia respectively.

3.2 Robust multiperiod poverty comparisons for the two-period case

In the following we first empirically illustrate how to test for robustness to poverty lines, and then show how to test for robustness to the aggregation procedure.

3.2.1 Robustness to poverty lines

To illustrate the robustness to poverty lines we use three waves of the Peruvian household panel data and consider the time spans 1998 to 1999 and 1999 to 2000. The overlap of the two time spans does not pose any problem for the test of robustness to poverty lines. However, it would not allow for a test of the robustness to the aggregation procedure since, as discussed above, in the case of an overlap, the same weight has to be attributed to each period.

According to Equation ??, poverty comparisons can be made by testing for significant differences between the dominance surface for 1998/99 and the dominance surface for 1999/2000. Testing robustness to the poverty line implies testing all points on the bisector between income in period 1 and income in period 2. Figure ?? shows the dominance surface of the first time span 1998-1999. The x and y axes measure income (or more precisely household expenditure per capita per day) at the beginning (1998) and the end (1999) of the time span. Expenditures are expressed in 2002 US\$ PPP equivalents. The third axis measures the cumulative share of people who are below the points defined in the (x, y) domain.

[insert Figure 2 about here]

Figure ?? shows the difference between the dominance surfaces of the time spans 1999/98 and 1999/2000. The relevant points can be found at the bisector of the graph, since we are testing only robustness to the poverty line (i.e., $z_1 = z_2$). The figure shows that for very low incomes, multiperiod poverty was higher in the first than in the second time span. However, as we increase the poverty line, we find that the cumulative share of people having had an income below that poverty line increases faster and that multiperiod poverty becomes higher for the second time span. This is a very interesting result because it already shows the importance of conducting dominance tests in this context. It can be seen even more clearly in Table ??.

[insert Figure 3 and Table 1 about here]

The vertical axis in Table ?? shows income at the beginning of the time spans and the horizontal axis at the end of the time spans. The value ‘1’ indicates a significant positive difference, i.e., 1999/2000 dominates 1998/99. ‘0’ means an insignificant difference, while ‘-1’ indicates a significant negative difference, i.e., 1998/99 dominates 1999/2000. Actually, we should check for poverty dominance at every possible point on this bisector, i.e., at every possible poverty line (e.g., \$1, \$1.01, \$1.02, etc.). However, for the sake of a simple and transparent presentation, we abstained from this detailed analysis and just checked at all poverty lines that were multiples of \$0.5. Again, the table demonstrates the relevance of our approach. Given

the 1\$ poverty line, one would conclude that ‘chronic’ poverty, i.e., individuals who are under the poverty line in both periods constituting a time span, would have fallen from the first to the second time span because there were more people with less than \$1PPP in 1998 *and* 1999 than in 1999 *and* 2000. However, if one were to choose the \$2 poverty line, it would not be possible to conclude poverty dominance because of the only insignificant differences. If, on the other hand, the \$3 poverty line were chosen, one would conclude a rise in chronic poverty from the first to the second time span. Thus, any conclusion about poverty orderings relies heavily on the poverty line chosen. In other words, to state that ‘chronic’ poverty (as defined here) has changed significantly from one time span to another, one should, in a first step, define an appropriate maximum poverty line and then check, in a second step, whether poverty dominance holds at every possible poverty line up to this maximum. Our methodology makes it possible to follow this procedure.

3.2.2 Robustness to aggregation procedures

Robustness to aggregation procedures seems important since the weights attributed to different periods are often arbitrary. ‘Time discounting’, for instance, might appear to be the most appropriate weighting scheme for economists. However, it is empirically very difficult to obtain a reliable and precise value for consumers’ discount rates. One therefore needs to be sure that the poverty ordering is robust against alternative weights in a reasonable range. Variations in the discount rate mean changes in the aggregation procedure across periods within a time span. It should be noted that when making poverty comparisons, changing the discount rate is equivalent to applying different poverty lines to period 1 and period 2. Put differently, increasing the poverty line in the second period is equivalent to applying a discount rate to the second period. As will be demonstrated now, if the time spans under consideration do not overlap, our methodology simultaneously ensures robustness to poverty lines *and* aggregation procedures (and in any case to a wide range of poverty measures).

We compare the time span 1998/1999 with the time span 2000/2001. In contrast to the procedure illustrated above, one now has to check not only for significant differences between the two surfaces at the bisector but at all points below and above the bisector up to a reasonable maximum poverty line. This becomes clear when looking at Figure ?? and Table ?. Figure ?? shows the difference between the two dominance surfaces. A robust poverty ordering would require that one surface be above the other surface at all points up to a reasonable maximum poverty line. This is obviously not the case here. Table ? substantiates the impression of missing dominance. Given the many ‘0’s’ in the grid of test points, it is clear that poverty dominance cannot be established for any reasonable set of poverty lines in any aggregation procedure.

[insert Figure 4 and Table 2 about here]

To underline the economic relevance of our approach, we now show the specific outcomes of weighting period 1 and period 2 differently. We consider poverty orderings $\mathbf{P}(\mathbf{Z}, \mathbf{R})$ which are robust for a broad class of poverty measures, $P(F; z; r)$ and a large range of poverty lines, $z \in Z$ and discount rates, $r \in R$. Hence, we rely on a poverty index P that assesses the degree of poverty, given a two-period distribution $F(y_1, y_2)$ when the poverty line is z and the discount factor of subsequent periods to the first period of a given time span is r . Therefore, we say that

$$F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z}, \mathbf{R})F(y_{1b}, y_{2b})$$

if and only if $P(F(y_{1a}, y_{2a}); z, r) \leq P(F(y_{1b}, y_{2b}); z, r) \forall z \in Z; r \in R$

and $P(F(y_{1a}, y_{2a}); z, r) < P(F(y_{1b}, y_{2b}); z, r)$ for some $z \in Z; r \in R$, (5)

where $F(y_{1a}, y_{2a})\mathbf{P}(\mathbf{Z})F(y_{1b}, y_{2b})$ means that multiperiod poverty is less over time span T_A than over time span T_B with respect to the poverty index P , the range Z and any weighting factor in the range R to discount incomes observed in later periods to the first period constituting a time span.

To illustrate this methodology, we consider the comparison of the time spans 1998/99 and 2000/2002. The two time spans have different lengths, such that discounting to the present is important. The results are shown in Figure ?? and Table ?. Table ? has two dimensions. The first dimension corresponds to income, and the second corresponds to the discount rate used. That means that each cell corresponds to one point of the bisector between income in the first and second periods of each time span, where income in the second period is discounted by the factor $(1 + r)^{-n}$, where n is the length of the respective time span measured in years. For instance, the ‘1’ in the sixth column of the first row means that if incomes of period 2 in each time span are discounted by a factor 1.05 per year, multiperiod poverty was significantly higher in 1998/99 than in 2000/01. Here, we only checked at a grid of test points for significant differences of the bisectors for a large range of discount rates and poverty lines. Overall, Table ? shows that in this comparison, 1998/99 vs. 2000/2002, poverty dominance does not hold for the whole grid. The ordering clearly depends on the discount rate and poverty line chosen. Theoretically, we could also test points below or above the bisector, i.e., setting different poverty lines in period 1 and 2 of each time span. However, from an economic point of view it makes sense to use a constant poverty line and feed any period-specific weighting into the discount rate.

[insert Figure 5 and Table 3 about here]

3.3 Robust multiperiod relative poverty comparisons for the two-period case within and across countries

We now apply our concept to relative poverty comparisons. The concept of multiperiod *relative* poverty comparison is closely related to the method of absolute poverty comparison introduced in Section ???. Absolute poverty measures deal with *income* mobility; they consider absolute poverty frontiers, for example the 1\$ PPP poverty line, and keep track of people who either stay below or cross this fixed frontier. Relative poverty measures, on the other hand, take into account *social* mobility; while still keeping track of people who either stay below or cross the poverty line, this frontier becomes endogenous, for example, expressed as a ratio of the median income. Embedding our concept of multiperiod poverty in the concept of relative poverty has some common features with Bossert’s, D’Ambrosio’s and Peragine’s (2006) concept of ‘social exclusion’.

To illustrate the idea of relative poverty, consider a household that has experienced a significant increase in income from one period to another and thus moved out of poverty from an absolute perspective. If the income of almost all households in the region has risen in a similar way, this household might still be poor from a relative perspective, i.e., the poverty gap to the median did not decline. Accordingly, people are referred to as ‘chronically poor’ in relative terms if their income, measured as a ratio of the median income, stays below a given proportion for consecutive years.

To test for differences in relative poverty between two time spans, we standardize household expenditures by a relative poverty line \tilde{z} , i.e., $\tilde{y} = y/\tilde{z}$. We choose $\tilde{z} = 50\%$ of median income.³ Accordingly, a relative income of 1, for example, means that the individual’s income is exactly half of the median income.

To illustrate the concept of relative multiperiod poverty, we first compare two time spans in Indonesia.

The difference in relative poverty between the time spans 1993/97 and 1997/2000 in Indonesia is presented in Figure ??. The x and y axes measure relative income, \tilde{y} , at the beginning and the end of the time spans. The figure does not show any systematic pattern. This is supported by Table ??, which shows the test points. Here the 0 in the third row of the third column, for example, means that the share of people who had less than 50% of the median income ($\tilde{y} = 1$) did not significantly change between the time spans 1993/97 and 1997/2000. Hence, no conclusions about changes in multiperiod poverty can be drawn.

[insert Figure 6 and Table 4 about here]

Our concept of relative poverty orderings is also applicable to cross-country comparisons. Absolute poverty comparisons using US\$ PPP equiv-

³Note that it does not matter which share of the median is used as poverty line.

alents as denominations of national currencies are interesting if countries have comparable and rather low living standards. But for countries with very different living standards or for very rich countries, i.e., those where absolute poverty is more or less negligible, standardizing expenditures in the way proposed in this section might be more informative. Here, one compares the extent of poverty relative to the country-specific average income to derive an impression of inequality between the poorest and the median income earner. To illustrate this, we now compare Peru to Indonesia. Peru has a median income of 4.7\$ PPP and Indonesia of 3.7\$ PPP per person per day. For these two countries, we consider the time span 1997/2000 with income observations in 1997 and 2000 for each.

Table ?? shows the matrix of test points of differences of the two-period poverty surfaces (‘Peru minus Indonesia’). Relative poverty seems to be higher in Peru. Even though dominance cannot be established over the entire domain, the maximum poverty set for relative dynamic poverty is wide enough to conclude dominance. The proportion of poor individuals is higher in Peru no matter what ‘reasonable’ relative poverty line or aggregation procedure is chosen.

[insert Table 5 about here]

3.4 Robust multiperiod poverty comparisons for the n -period case

As a matter of course, poverty comparisons over two time spans demand panel data over multiple periods. Consequently, the question arises how the time spans under consideration should be constructed if more than two periods are available within each. Which period should be the end of the first and the beginning of the second time span? How many periods should constitute a time span? These are very general questions regarding the measurement of multiperiod poverty (or chronic poverty more specifically). Depending on the panel data available, often several different time span constructions are possible, varying in time span length and the number of periods taken into account. This raises the question, for example, whether comparisons should be made with the maximum overlap (e.g., $T_A[y_1, y_2, \dots, y_{n-1}]$ vs. $T_B[y_2, y_3, \dots, y_n]$), without any overlap ((e.g., $T_A[y_1, y_2, \dots, y_{n/2}]$ vs. $T_B[y_{n/2+1}, y_{n/2+2}, \dots, y_n]$), or with something in between. Depending on these choices, poverty orderings may differ. Thus, beyond robustness to poverty indices, poverty lines and aggregation procedures, one may also require poverty comparisons to be robust to the construction of the time spans.

To illustrate this, we use five waves of the Peruvian household panel data (1998-2002). To simplify matters, we require that in each comparison, the first period of time span T_A be 1998 and the last period of time span

T_B be 2002. We also abstain from making comparisons for different time span lengths. However, the rest of the time span construction is arbitrary and consequently, any poverty ordering may depend on how exactly the construction is carried out. There are five different comparisons that make sense from an economic point of view: three where we consider time spans comprising two periods, one where we consider time spans comprising three periods, and one where we consider time spans comprising four periods:

$$\begin{aligned}
 & [1998; 2000] \text{ vs. } [2000; 2002] \\
 & [1998; 1999] \text{ vs. } [2001; 2002] \\
 & [1998; 2001] \text{ vs. } [1999; 2002] \\
 & [1998; 1999; 2000] \text{ vs. } [2000; 2001; 2002] \\
 & [1998; 1999; 2000; 2001] \text{ vs. } [1999; 2000; 2001; 2002]
 \end{aligned}$$

Given the difficulty in determining which of these five comparisons is most appropriate, dynamic poverty comparisons should be robust to all of them. For example, one can imagine a case in which the \$3 poverty line is considered to be a reasonable maximum poverty line when comparing poverty dynamics for the time span 1998 - 2002 in Peru. In this case, the poverty ordering is only considered robust if poverty dominance can be established for every possible poverty line up to the \$3 poverty line *and* for every above-mentioned type of construction for the time spans. In the following, we do not consider robustness to the aggregation procedure.

Table ?? shows the results of such a dominance test. Obviously, according to our proposed methodology, no significant ordering of poverty dynamics can be established for the time span 1998 to 2002. This is a very interesting result given the large number of 1's in Table ?. Imagine you wanted to assess the development of chronic poverty for the time span 1998 to 2002: If you used the \$2 poverty line and compared the time spans [1998; 1999; 2000] and [2000; 2001; 2002]—which might be judged a reasonable comparison at first glance—you would spuriously have concluded that chronic poverty has fallen. However, taking the time spans [1998; 1999] and [2001; 2002] shows instead that no conclusion can be drawn. Again, the poverty ordering depends heavily on the chosen poverty line and the way the time spans are constructed.

[insert Table 6 about here]

4 Discussion

In this paper, we presented a concept allowing to undertake multiperiod poverty comparisons over time and space without arbitrarily aggregating income over various years. Inspired by the multidimensional stochastic dominance methodology elaborated by Duclos *et al.* (2006a), we created

n -period income surfaces for different time spans. These surfaces were then ordered using dominance tests. Once dominance is established, the poverty ordering is robust to a wide range of poverty indices, to a wide range of poverty lines, and to a wide range of aggregation procedures.

Furthermore, we extended our framework to the measurement of *relative* poverty. In contrast to the analysis of the dynamics of absolute poverty, which deals with income mobility, the analysis of the dynamics of relative poverty deals with social mobility. Given the political relevance of social mobility and inequality between the poorest and the median income earner, we think these measures offer an interesting contribution to the existing literature on poverty dynamics.

To illustrate our methodology, we compared poverty across time spans in Peru and between Peru and Indonesia. Furthermore, we highlighted the general problem of dynamic poverty comparisons, i.e., how time spans should be constructed. Which period should be the end of the first and the beginning of the second time span? How many periods should constitute a time span? We dealt with these questions in a demonstration of how to test robustness with respect to various construction modes.

However, the approach suggested and the ideas developed in this paper also have their limitations. The most important of these is certainly that all results are based on a sample of expenditures declared by households and that these declarations are generally affected by measurement error, which affects the bivariate distribution $F(y_1, y_2)$ (and n -variate distribution) much more than the univariate distribution $F(y)$. In fact, many empirical studies show that measurement error is such that the extent of β -convergence over time is overestimated (see e.g., Bound, Brown and Mathiowetz (2001); Breen and Moisisio (2004); Grimm (2006)). For our case, this would imply that multiperiod poverty is underestimated. In the absence of information on ‘true income’ or any instruments, there is not much that can be done about this, but it should be kept in mind when interpreting our results. However, the problem is obviously not specific to our approach but inherent in most approaches to the analysis of poverty dynamics.

References

- Breen, R. and P. Moisisio (2004), Poverty dynamics corrected for measurement error. *Journal of Economic Inequality*, 2 (3): 171 - 191.
- 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.
- Bossert W., C. D’Ambrosio and V. Peragine (2006), Deprivation and Social Exclusion. *Economica*, forthcoming.

- Bourguignon F. and S.R. Chakravarty (2002), Multidimensional poverty orderings. Delta Working Paper 2002-22, Delta, Paris.
- Bourguignon F. and S.R. Chakravarty (2003), The measurement of multidimensional poverty. *Journal of Economic Inequality*, 1: 25-49.
- Chronic Poverty Research Center (2004), *The Chronic Poverty Report 2004-05*. Institute for Development Policy and Management, University of Manchester.
- Cruces G. (2005), Income Fluctuations, Poverty and Well-Being Over Time: Theory and Application to Argentina. Discussion Paper No. DARP 76, London School of Economics.
- Duclos J.-Y., A. Araar and J. Giles (2006), Chronic and Transient Poverty: Measurement and Estimation, with Evidence from China. IZA Discussion Paper No. 2078, IZA, Bonn.
- Duclos J.-Y., D. Sahn and S.D. Younger (2006a), Robust Multidimensional Poverty Comparisons. *Economic Journal*, 116 (514): 943-968.
- Duclos J.-Y., D. Sahn and S.D. Younger (2006b), Robust Multidimensional Spatial Poverty Comparisons in Ghana, Madagascar, and Uganda. *World Bank Economic Review*, 20 (1): 91-113.
- Foster J.E., J. Greer and E. Thorbecke (1984), A Class of Decomposable Poverty Measures. *Econometrica*, 52: 761-776.
- Foster J.E. and A.F. Shorrocks (1988a), Poverty Orderings. *Econometrica*, 56: 173-177.
- Foster J.E. and A.F. Shorrocks (1988b), Poverty Orderings and Welfare Dominance. *Social Choice and Welfare*, 5: 179-189.
- Grimm M. (2007), Removing the anonymity axiom in assessing pro-poor growth. *Journal of Economic Inequality*, 5 (2): 179-197.
- Grootaert C. and R. Kanbur (1995), The Lucky Few amidst Economic Decline: Distributional Change in Côte d'Ivoire as seen through Panel Datasets, 1985-1988. *Journal of Development Studies*, 31 (4): 603-619.
- Heston, A., R. Summers and B. Aten (2002), Penn World Table Version 6.1, Center for International Comparisons at the University of Pennsylvania (CICUP).
- Hulme D. and A. McKay (2005), Identifying and Measuring Chronic Poverty: Beyond Monetary Measures. Mimeo, Institute for Development Policy and Management, University of Manchester.

- Hulme D. and A. Shepard (2003), Conceptualizing Chronic Poverty. *World Development*, 31 (3): 403-424.
- Jalan J. and M. Ravallion (1998), Transient Poverty in Postreform Rural China. *Journal of Comparative Economics*, 26: 338-357.
- 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

Figure 1: Test domain for dynamic poverty comparisons

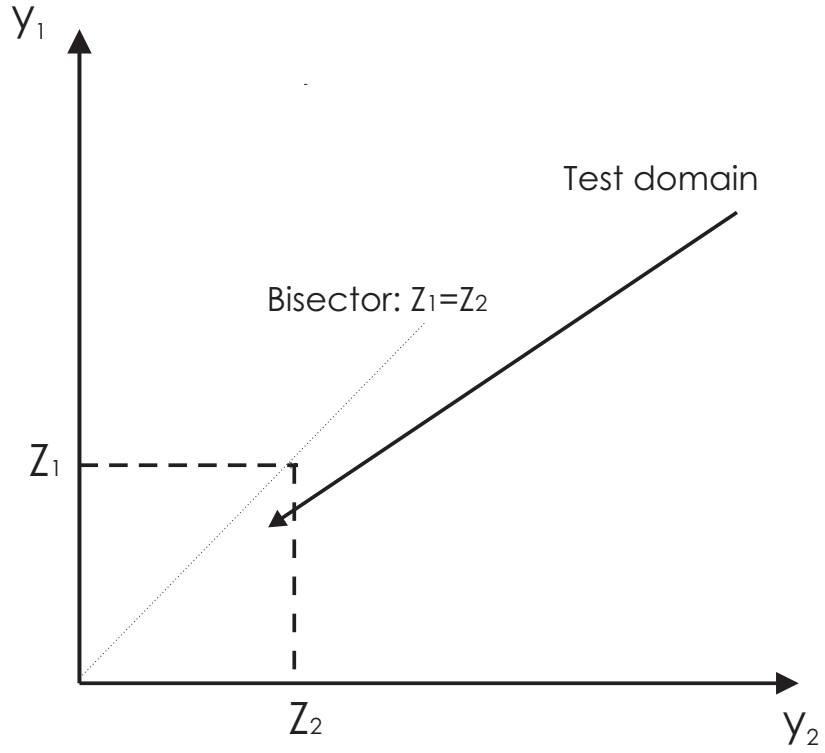


Figure 2: Poverty in Peru: Dominance surface of the time span 1998/99




Figure 3: Poverty in Peru: Differences in dominance surfaces (1998/99 - 1999/2000)




Figure 4: Poverty in Peru: Differences in dominance surfaces (1998/99 - 2000/01)




Figure 5: Poverty in Peru: Differences in dominance surfaces (1998/99 - 2000/02)




Figure 6: Relative poverty in Indonesia: Differences in dominance surfaces (1993/97 - 1997/2000)




Table 1: Poverty in Peru
Differences in dominance surfaces (1998/99 - 1999/2000)

Income period 1	Income period 2									
	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0	
1.0	1									
1.5		0								
2.0			0							
2.5				0						
3.0					-1					
3.5						-1				
4.0							-1			
4.5								-1		
5.0									-1	

Income: Household income per capita;
1 indicates that the 1998/99 surface was significantly above the 1999/2000 surface, -1 indicates the opposite, 0 indicates no significant difference.

Table 2: Poverty in Peru
Differences in dominance surfaces (1998/99 - 2000/01)

Income period 1	Income period 2								
	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
1.0	0	0	0	0	0	0	0	0	0
1.5	0	0	0	0	0	0	0	0	0
2.0	-1	0	0	0	0	0	0	0	0
2.5	0	0	0	0	0	0	0	0	0
3.0	0	0	0	0	0	0	-1	-1	-1
3.5	0	0	0	0	0	-1	-1	-1	0
4.0	0	-1	-1	0	-1	-1	-1	-1	-1
4.5	0	-1	-1	0	-1	-1	-1	-1	-1
5.0	0	-1	-1	-1	-1	-1	-1	-1	-1

Income: Household income per capita;
1 indicates that the 1998/99 surface was significantly above the 2000/01 surface, -1 indicates the opposite, 0 indicates no significant difference.

Table 3: Poverty in Peru
Differences in dominance surfaces (1998/99 - 2000/02)

Income	Discount rate of period 2															
	.00	.01	.02	.03	.04	.05	.06	.07	.08	.09	.1	.11	.12	.13	.14	.15
1.0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
1.2	1	1	1	1	1	1	0	0	0	0	0	0	0	0	0	1
1.4	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1
1.6	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1
1.8	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2.0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2.2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2.4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2.6	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
2.8	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
3.0	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
3.2	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
3.4	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1
3.6	0	0	0	0	0	1	1	1	1	1	1	1	1	1	1	1
3.8	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1
4.0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1
4.2	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	1
4.4	0	0	0	0	0	0	0	0	0	1	1	1	1	1	1	1
4.6	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1
4.8	0	0	0	0	0	0	0	1	1	1	1	1	1	1	1	1

Income: Household income per capita;
1 indicates that the 1998/99 surface was significantly above 2000/02 surface, -1 indicates the opposite, 0 indicates no significant difference.

Table 4: Relative poverty in Indonesia
Differences in dominance surfaces: (1993/97 - 1997/2000)

\tilde{Income} period 1	\tilde{Income} period 2													
	0.8	0.9	1.0	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0	
0.8	1	1	0	0	0	0	0	1	1	1	1	1	1	
0.9	1	0	0	0	0	0	0	1	1	1	1	0	0	
1.0	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.1	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.2	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.3	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.4	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.5	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.6	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.7	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.8	0	0	0	0	0	0	0	0	0	0	0	0	0	
1.9	0	0	0	0	0	0	0	0	0	0	0	0	0	
2.0	0	0	0	0	0	0	0	0	0	0	0	0	0	

\tilde{Income} is household income per capita, standardized by a relative poverty line, $\tilde{z} = 50\%$, of median income: $\tilde{Income} = Income/\tilde{z}$
1 indicates that the 1993/97 surface was significantly above the 1997/2000 surface, -1 indicates the opposite, 0 indicates no significant difference.

Table 5: Relative poverty in Peru and Indonesia
Differences in dominance surfaces: Peru (1997/2000) - Indonesia (1997/2000)

\tilde{Income} period 1	\tilde{Income} period 2													
	0.8	0.9	1.0	1.1	1.2	1.3	1.4	1.5	1.6	1.7	1.8	1.9	2.0	
0.8	1	1	1	1	1	1	1	1	1	1	1	1	1	
0.9	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.0	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.1	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.2	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.3	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.4	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.5	1	1	1	1	1	1	1	1	1	1	1	1	1	
1.6	1	1	1	1	1	1	1	1	1	1	1	0	1	
1.7	1	1	1	1	1	1	1	1	1	0	0	0	0	
1.8	1	1	1	1	1	1	1	1	0	0	0	0	0	
1.9	1	1	1	1	1	0	0	0	0	0	0	0	0	
2.0	1	1	1	1	1	0	0	0	0	0	0	0	0	

\tilde{Income} is household income per capita, standardized by a relative poverty line, $\tilde{z} = 50\%$, of median income: $\tilde{Income} = Income/\tilde{z}$
1 indicates that the Peru surface was significantly above the Indonesia surface, -1 indicates the opposite, 0 indicates no significant difference.

Table 6: Poverty in Peru
Differences in dominance surfaces for several construction modes of time spans

Income	[98; 00] vs. [00; 02]	[98; 99] vs. [01; 02]	[98; 01] vs. [99; 02]	[98; 99; 00] vs. [00; 01; 02]	[98; 99; 00; 01] vs. [99; 00; 01; 02]
1	0	0	0	0	0
1.5	0	0	0	0	0
2	0	0	1	1	1
2.5	1	0	1	1	1
3	1	0	0	0	1
3.5	1	0	0	0	0
4	0	0	0	0	0
4.5	0	0	0	0	0
5	0	0	0	0	0

Income: Household income per capita;
1 indicates that the earlier surface was significantly above later surface, -1 indicates the opposite,
0 indicates no significant difference.