The Econometrics of Income and Social Mobility

The Growth Incidence Curve and Related Questions

Michel Lubrano

January 2019

Contents

1 Introduction

This is a first approach to income dynamics which studies the evolution of the income distribution between two periods of time. Illustrated by the Elephant curve:

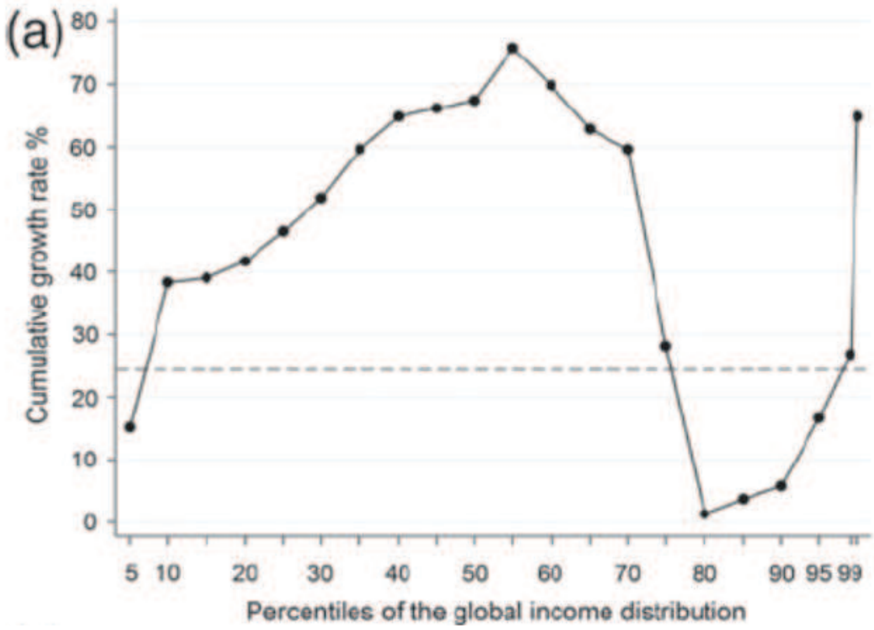

Figure 1: The Elephant Curves of Lakner and Milanovic (2016)

What is present on this curve? We have

- on the ordinates the growth rate in percentage of the world economy and
- on the abscissae the quantiles of the World Income Distribution.

The question now is to know how this curve is built. It is clearly related to income dynamics, because it indicates which quantiles of the income distribution benefited the most of the economic growth. If this curve were flat, everybody would have benefited equally. This is not the case in the above Figure.

Several research questions were asked around this curve

1. The first question was: is growth pro-poor or not? The initial paper was Ravallion and Chen (2003). Put in another terms: is the trickle-down theory verified or not? The Growth Incidence Curve (GIC) of Ravallion and Chen (2003) is one way of comparing changes in an income distribution between two points in time. It measures how growth is distributed over the quantiles.

2. Is there another way of building such curves? They are related to the Generalized Lorenz curve. The approach of Son (2004) with the Poverty Growth Curve (PGC).

3. When comparing two curves, how is it related to stochastic dominance? Duclos (2009) and Araar et al. (2009).

4. Do these curves really represent dynamics? Bourguignon (2011) introduced the non-anonymous growth incidence curve, starting this time from a joint bivariate distribution and deriving a modified growth incidence curve, having different properties. The problem is that it is quite difficult to define bivariate quantiles.

2 Quantiles and Lorenz curves

In order to introduce the topic, we have first to detail some concepts, in particular those related to the quantile function, which will be helpful in order to redefine in an uncommon way some well-known results.

2.1 CDF and quantiles

Let Y (Household Income divided by an equivalence scale, GDP per capita, ...) be a continuous random variable with cumulative distribution function (cdf) $F(y)$ and probability density function (pdf) $f(y)$ with support contained on the non-negative real line. The quantile function is defined as the inverse of the cdf:

$$Q(p) = F^{-1}(y), \qquad \text{or} \qquad Q(p) = \inf_{y \geq 0}(F(y) \geq p).$$

Remark:

> **How to estimate the cdf?** The empirical income distribution is formed by n observations of Y, noted y and arranged by increasing order. The sequence of order statistics is noted $[y_{[i]}]$. The graph of the empirical quantile function is obtained by plotting the n component vector $[p_i = i/n]$ in $[0, 1]$ against the n order statistics. If we normalize this graph by the mean, we get the well-known Pen's parade.

For computing a mean, a poverty index, an inequality index or the Lorenz curve, we can use directly the pdf of the random variable Y. For instance the mean is defined as:

$$\mu = \int_0^\infty y f(y)\, dy. \tag{1}$$

However, we can use a dual estimator based on the quantile function. Let us consider the change of variable $y = F^{-1}(p)$ and apply it to (1), we get:

$$\mu = \int_0^1 F^{-1}(p)\, dp = \int_0^1 Q(p)\, dp. \tag{2}$$

If z is a poverty line (to be defined later on), then the FGT poverty index introduced by Foster et al. (1984) is defined as:

$$P_\alpha(y, z) = \int_0^z (1 - y/z)^\alpha f(y)\, dy.$$

Using the same change of variable, we have

$$P_\alpha(y, z) = \int_0^q (1 - Q(p)/z)^\alpha \, dp,$$

with $q = F(z)$.

2.2 The Lorenz curve

Let us now consider the Lorenz curve, a widely used graphical representation of inequality introduced in Lorenz (1905). It was originally defined by two equations:

$$L(p) = \frac{1}{\mu} \int_0^y t\, f(t)\, dt \tag{3}$$
$$p = F(y). \tag{4}$$

Using the same change of variable $y = Q(p)$, Gastwirth (1971) provides the following form of the Lorenz curve:

$$L(p) = \frac{\int_0^p Q(t)\, dt}{\int_0^1 Q(t)\, dt},$$

which immediately relates the quantile function to the Lorenz curve as:

$$\mu L(p) = \int_0^p Q(t)\, dt \Rightarrow Q(p) = \mu L'(p). \tag{5}$$

From this expression it becomes clear that the mean income in the population is found at the percentile at which the slope of $L(p)$ (i.e. $L'(p)$) is equal to 1.

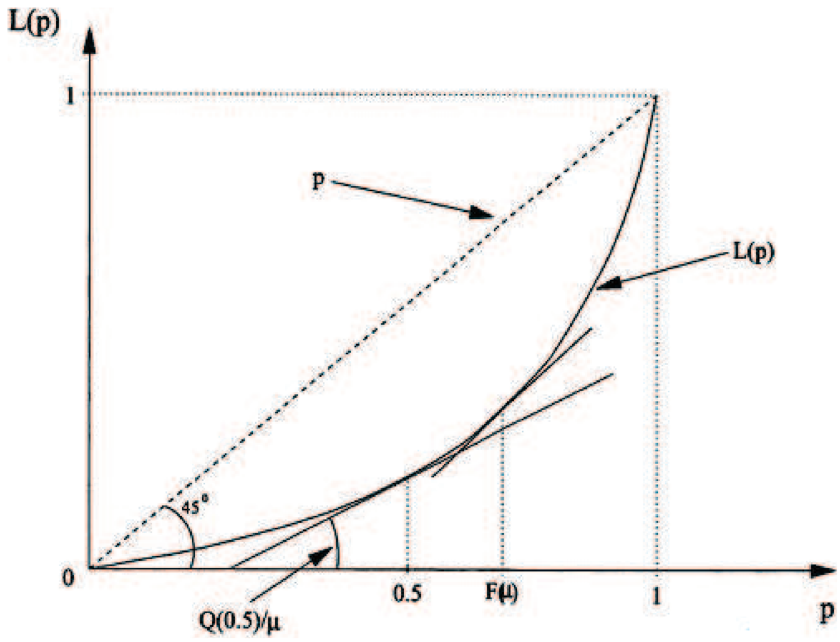

Figure 2: Lorenz curve (source Duclos and Araar 2006)

Remark:

The Lorenz curve is easy to estimate in a distribution free approach. Let us order the observations so as to define the order statistics $y_{[i]}$.

$$L(p = i/n) = \frac{1}{\bar{y}} \sum_{j=1}^{i} y_{[j]}$$

This estimator can be quite irregular if there are not many observations, so it can be useful to adjust a parametric distribution to the income data and consider the corresponding parametric form of the Lorenz curve. For instance, if y is modelled according to a Lognormal distribution $f_\Lambda(y|\mu, \sigma^2)$, the corresponding Lorenz curve is:

$$L(p) = \Phi(\Phi^{-1}(p) - \sigma),$$

while the Gini index is:

$$I_G = 2\Phi(\sigma/\sqrt{2}) - 1.$$

Both quantities depend only on the parameter σ which can be estimated as the standard deviation of the log of the observations.

2.3 The Generalized Lorenz Curve

A welfare function can be expressed as

$$W(y) = \mu(1 - I_y),$$

where I_y represents an inequality index at value between 0 and 1. The Lorenz curve is a representation of inequality. For instance the Gini index is twice the area between the 45° line and the Lorenz curve. So the Lorenz curve represents only one part of welfare as the mean is missing. A variant of the Lorenz curve representing inequality while taking into account the level of income has been introduced formally in Shorrocks (1983), that is the generalized Lorenz curve. It is simply obtained by multiplying the Lorenz curve by the mean income μ:

$$GL(p) = \int_0^p Q(t)\, dt = \mu L(p).$$

When the Lorenz curve was contained in the unit square, the Generalized Lorenz curve is contained in the rectangle with a 0-1 base and a height covering $0 - \mu$.

Numerous inequality and poverty measures also rely on the quantile function, and thus, can be derived from the Lorenz curve too. This is illustrated for instance in Foster and Shorrocks (1988).

3 Inequality dynamics

These preliminaries having been now detailed, we turn to the proper definition of the GIC. Let us consider two dates t and $t-1$ and their respective distributions $F_t(y)$ and $F_{t-1}(y)$. We shall define two types of curves: The GIC with Ravallion and Chen (2003) and the PGC with Son (2004). Both articles have the term *pro-poor growth* in their title.

3.1 The growth incidence curve GIC

The growth incidence curve, introduced in Ravallion and Chen (2003), measures the growth rate of the p-quantile for every p:

$$g_t(p) = \frac{Q_t(p)}{Q_{t-1}(p)} - 1 \simeq \log Q_t(p) - \log Q_{t-1}(p). \tag{6}$$

Using (5), the GIC can be immediately related to the Lorenz curve with:

$$g_t(p) = \frac{L'_t(p)}{L'_{t-1}(p)}(\gamma_t + 1) - 1 \simeq \log GL'_t(p) - \log GL'_{t-1}(p), \tag{7}$$

where $\gamma_t = (\bar{y}_t - \bar{y}_{t-1})/\bar{y}_{t-1} \simeq \log(\bar{y}_t) - \log(\bar{y}_{t-1})$ is the average growth rate. Two immediate properties can be derived from (7):

1. if inequality does not change then $g_t(p) = \gamma_t$ for all p,

2. the p-quantile increases if $g_t(p) > 0$.

Thus the GIC corresponds to the variation of the first derivative of the generalized Lorenz curve. Graphically, the GIC associates the growth rate of income with respect to proportion p of individuals ordered by increasing income. By drawing the horizontal line corresponding to the rate of growth of the mean income (or the median income), the quantiles below that line have a rate of growth of their income which is lower than the growth rate of the mean income (or the median income). This is the anonymous growth incidence curve. We do not know if and we do not impose that the same household or individual is followed over time. We consider only a representative quantile, without knowing the identity of its participants. *Consequently, we do not need a balanced panel to make inference.*

Remark:

How to estimate the GIC? A distribution free approach is easy to implement as it is related to estimating the quantile function:

$$GIC(p) = log(\hat{Q}_t(p)) - \log(\hat{Q}_{t-1}(p)).$$

For instance is `R`, that would give

```
gt_np = function(p,y1,y2){
    g2 = quantile(y2,p)
    g1 = quantile(y1,p)
    g = log(g2)-log(g1)
    g
}
```

However, in the case of a small sample, very few point would be devoted to estimating a particular quantile. So the method is not very precise. In this case a parametric assumption might be valuable. For the lognormal distribution, we would have:

$$g_t(p) = (\mu_t + \sigma_t^2 \Phi^{-1}(p)) - (\mu_{t-1} + \sigma_{t-1}^2 \Phi^{-1}(p)).$$

However, this form is very much constrained by the shape of $\Phi^{-1}(p)$. So that a mixture assumption for the income distribution would be preferred. However, in this case, the quantile function for the mixture is much more difficult to estimate because there is no analytical form for the quantile function of a mixture. The CDF of the mixture has to be inverted numerically.

3.2 The Poverty Growth Curve

An alternative approach for assessing distributional changes has been proposed by Son (2004) who introduces the poverty growth curve (PGC). The initial question of Son (2004) was to determine whether the mean income of the lower quantiles (corresponding to the poor) is growing quicker than the mean income of the other quantiles. The PG curve is defined as the variation in percentage of the average income of the bottom $p\%$ of the population and corresponds to $\Delta \log(\bar{y}_p)$, where $\bar{y}_p$ is the average income of the bottom $p\%$. Because, using (2) the Lorenz curve can be written as:

$$L(p) = \frac{\int_0^p Q(t)\,dt}{\int_0^1 Q(t)\,dt} = \frac{p\bar{y}_p}{\bar{y}},$$

the PG curve corresponds to:

$$PGC_t(p) = \Delta \log \bar{y}_p = \Delta \log \bar{y} + \Delta \log L(p) = \Delta \log GL(p). \tag{8}$$

In this context, growth is pro-poor if the variations of the Lorenz curve are positive for all p up to a given value. The poverty growth curve (PGC) is thus equal to the variation of the generalized Lorenz curve. Because $L(p = 1) = 1.0$ then $\Delta L(p = 1) = 0$ and the PGC is equal to γ_t at $p = 1$.

1. Growth is pro-poor when $PGC_t(p) > \gamma_t$, which means that the $PGC_t(p)$ is decreasing in p as $PGC_t(p = 1) = \gamma_t$.

2. Poverty simply decreases when $PGC_t(p) > 0$ for all $p < 1$.

3. When $0 < PGC_t(p) < \gamma_t$ for all $p < 1$, there is a phenomenon of *trickle-down* growth, that is to say poverty is reduced but not as much as it could because the rich are receiving proportionally more.

Using the same notation and approximation as before, the GIC can be written as:

$$g_t(p) = \Delta \log Q(p) = \Delta \log \bar{y} + \Delta \log L'(p) = \Delta \log GL'(p),$$

so that the two curves can be compared. Both measures are obtained as the variation of the log of the mean income plus the variation of either the log of the Lorenz curve or the log of its derivative. While the *growth rate of income at the p-quantile* is used for the GIC, the PGC is based on the estimation of the *growth rate of the mean income up to the p-quantile.*

Remark:

How to estimate the PGC? We can use the same distribution free assumption as for the GIC. We have a difficulty as the value of p has to be the same for the two generalized Lorenz curves. A way to obtain this can be in `R`:

```
Gt_np = function(p,y1,y2){
    m1 = mean(y1)
    m2 = mean(y2)
    L1 = (cumsum(sort(y1))/length(y1)/m1)[p*length(y1)]
    L2 = (cumsum(sort(y2))/length(y2)/m2)[p*length(y2)]
    le = log(m2) - log(m1)
    G = le+log(L2/L1)
    G
}
```

But simpler solutions are certainly possible, using the `ineq` package and its function for estimating a Lorenz curve. Under a lognormal assumption, one would have:

$$PGC(p) = (\mu_t + \sigma_t^2/2) - (\mu_{t-1} + \sigma_{t-1}^2/2) + \log\left(\frac{\Phi(\Phi^{-1}(p) - \sigma_t)}{\Phi(\Phi^{-1}(p) - \sigma_{t-1})}\right),$$

using the expressions for the mean and the Lorenz curve of the lognormal distribution.

3.3 Poverty lines and pro-poor growth

We have spoken about pro-poor growth, but how to define a poor and what is a poverty line? A poor household or individual has an income which is under the poverty line. The latter can be defined in several ways:

1. An absolute poverty line is based on the definition of the value of a basket of goods that are necessary to survive. It can be based on the cost of 2400 calories plus the cost of other basic services.

2. The international (extreme) poverty line was fixed at \$1.25, and now at \$1.90 because the defining of PPP has changed over time.

3. A relative poverty line, as adopted by the EU: 50% of the mean income or 60% of the median income

4. A subjective poverty line, based on surveys: the minimum income question.

So if the poverty line evolves with the mean income, there can be no gain in term of number of poor while growth was pro-poor. The GIC, by considering just quantiles answers to this question. But does it do it in a consistent way, obeying some axioms? We shall try to relate the GIC curve to a poverty index which verifies the usual axioms, the Watts poverty index. This index writes as:

$$P_W = \int_0^z \log\left(\frac{z}{y}\right) f(y)\, dy$$

By a change of variable $q = F^{-1}(z)$ where z is the poverty line and q the poverty headcount, we get

$$P_W = \int_0^q \log\left(\frac{z}{Q(p)}\right) dp.$$

Ravallion and Chen (2003) suppose that z is invariant over time, while y is varying over time. So by differentiating with respect to time, they get

$$-\frac{dP_W}{dt} = \int_0^q \frac{d\log(Q(p))}{dt}\, dp = \int_0^q GIC(p)\, dp.$$

So the area under the GIC up to the headcount ratio gives minus the change in the Watts index.

Compared to the GIC, the PGC seems more robust as

1. The GIC is based on the growth rate of per capita income at the p^{th} percentile. It is estimated using unit record data (individual data). Depends on the accuracy of the Lorenz curve estimates.

2. PGC is based on the growth rate of the mean income up to the p^{th} percentile. Only decile or quintile shares and mean income are required.

4 An empirical illustration using UK data

As an empirical illustration, we are going to use the Family expenditure survey , data collected for four years 1979, 1988, 1992 and 1996. We consider disposable income, equivalised by the McClements adult-equivalence scale and deflated by the corresponding relative consumer price index. This span of years cover the period when Margaret Thatcher was in power and when many changes occurred in the income distribution. After that period, for instance the value of the Gini coefficient remained stable as underlined in Jenkins (2000).

4.1 Stylized facts

The period when Margaret Thatcher was in power (1979-1988) is characterized by a large increase in both inequality and poverty. The next period (1992-1996) shows a slight decrease for poverty and inequality, but without returning to the low levels of 1979.

Table 1: Poverty and inequality in the UK

Year	1979	1988	1992	1996
Poverty	0.094	0.171	0.193	0.142
Theil	0.107	0.162	0.179	0.151

4.2 Growth incidence curves

A log linear model is implemented with heteroskedastic errors for each of the four sample dates. The implied GIC and PGC for 1979-1988 and for 1992-1996 are derived, using $np = 100$ points for the grid for p and $x = \bar{x}$, for having a general portrait and compare our estimates with those obtained by the traditional distribution-free approach. Figures 3 and 4 show that we have two opposed situations with well marked characteristics.

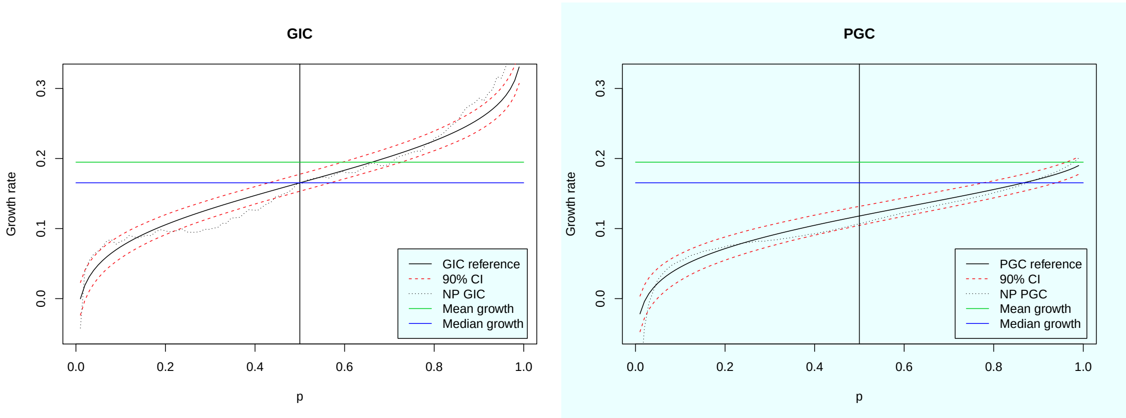

Figure 3: GIC and PGC with confidence intervals for 1979-1988

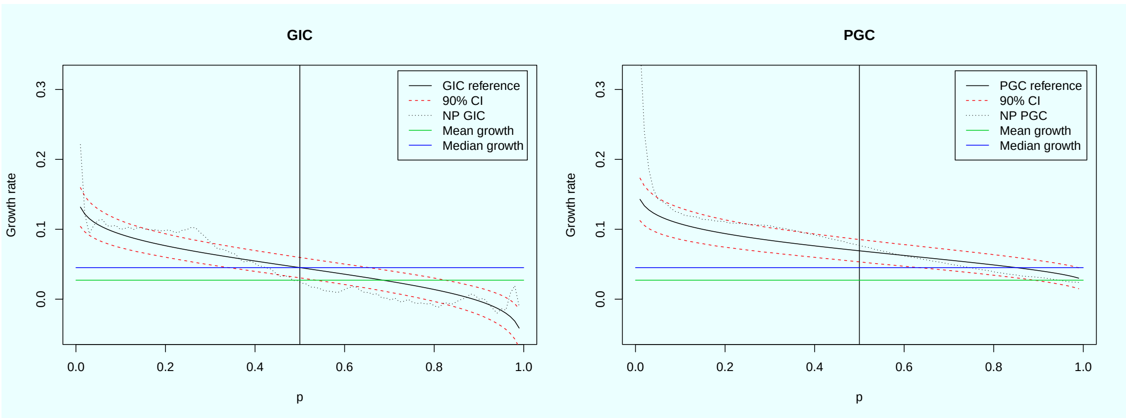

Figure 4: GIC and PGC with confidence intervals for 1992-1996

The first period 1979-1988 is clearly characterized by anti-poor growth while the second period 1992-1996 becomes pro-poor. For the first period, we must go up to the 0.65^{th} quantile in order to have an average growth rate of income greater than the average growth. During the second period, households up to the 0.65^{th} quantile experienced an income increase greater than average. This is not necessary to go to the second order with the PGC curves to confirm this conclusion. A striking characteristics is that once we adopt comparable scales the period is strongly anti-poor when the second period is mildly pro-poor.

We note finally that during the first period the distribution-free estimate are relatively closed to the parametric estimate and contained in a 90% posterior confidence interval. During the second period, the distribution-free estimate fluctuates more and does follow less the general pattern of our parametric estimate. This means that our model does not fit so well over the second period. As a matter of fact, the BICs are lower in absolute value for the second period than for the first one.[1]

5 Stochastic dominance

We shall now concentrate on showing how these two curves can be merged into the framework of stochastic dominance, following Davidson and Duclos (2000), Son (2004), Duclos (2009) and Araar et al. (2009).

Let us consider an income distribution growing between two periods with vectors of observation y_1 and y_2, a growth rate γ of y_2 over y_1 and a common poverty line z.

Definition 1. *Let us consider two income distributions y_1 and y_2 with respective CDF F_1 and F_2. We have stochastic dominance at the order one of y_2 over y_1 iff:*

$$F_2(z) < F_1(z) \qquad \forall z \in [0, \infty[.$$

Davidson and Duclos (2000) call the p-approach to dominance the comparison of the quantile functions of y_1 and y_2 noted respectively Q_1 and Q_2. Stochastic dominance at the order one of y_2 over y_1 corresponds to

$$Q_2(p) > Q_1(p) \qquad \forall p \in [0, 1].$$

The GIC compares in the same way two quantiles functions as shown in (6). Consequently, stochastic dominance of y_2 over y_1 at the first order is equivalent to:

$$g_t(p) > 0, \quad \forall p \in [0, 1].$$

Duclos (2009) and Araar et al. (2009) go a step further on and say that growth is relatively pro-poor if:

$$g_t(p) > \gamma_t, \quad \forall p \in [0, F_1(z)],$$

where $F_1(z)$ is the cumulative distribution of y_1 evaluated at z. This condition checks if the quantiles of the poor increase at a pace greater than average growth.

[1] In a linear regression model the BIC is equal to $n \log(1 - \mathrm{R}^2) + k \log n$ and so can be related to a measure of fit.

Stochastic dominance at the order one implies stochastic dominance at higher orders. The reverse is wrong. So it is easier to verify stochastic dominance at the order two than at the order one. Stochastic dominance at the order two is obtained by comparing generalized Lorenz curves instead of quantile functions.

Definition 2. *When $GL_2(p) > GL_1(p)$, $\forall p < 1$, we have second order stochastic dominance at the order two of y_2 over y_1.*

The PGC of Son (2004) compares also two Generalized Lorenz curves. Consequently the condition:

$$G_t(p) > 0, \quad \forall p \in [0, 1[,$$

is equivalent to testing for second order dominance. Duclos (2009) and Araar et al. (2009) go one step ahead and propose to test for:

$$G_t(p) > \gamma_t, \quad \forall p \in [0, \tilde{q}],$$

where

$$\tilde{q} = F_2[(1 + \gamma)z].$$

This corresponds to a situation where growth is said to be pro-poor by Son (2004). We see from (8) that is corresponds to a situation where the Lorenz curve in period 2 is entirely closer to the line of perfect equality.

References

Araar, A., Duclos, J.-Y., Audet, M., and Makdissi, P. (2009). Testing for pro-poorness of growth, with an application to Mexico. *Review of Income and Wealth*, 55(4):853–881.

Bourguignon, F. (2011). Non-anonymous growth incidence curves, income mobility and social welfare dominance. *The Journal of Economic Inequality*, 9(4):605–627.

Davidson, R. and Duclos, J.-Y. (2000). Statistical inference for stochastic dominance and for the measurement of poverty and inequality. *Econometrica*, 68:1435–1464.

Duclos, J.-Y. (2009). What is Pro-Poor? *Social Choice and Welfare*, 32(1):37–58.

Foster, J., Greer, J., and Thorbecke, E. (1984). A class of decomposable poverty measures. *Econometrica*, 52(3):761–766.

Foster, J. and Shorrocks, A. (1988). Poverty orderings. *Econometrica*, 56(1):173–177.

Gastwirth, J. (1971). A general definition of the Lorenz curve. *Econometrica*, 39(6):1037–1039.

Jenkins, S. P. (2000). Modelling household income dynamics. *Journal of Population Economics*, 13:529–567.

Lakner, C. and Milanovic, B. (2016). Global income distribution: From the fall of the Berlinwall to the great recession. *World Bank Economic Review*, 30(2):203–232.

Lorenz, M. O. (1905). Methods of measuring the concentration of wealth. *Publications of the American Statistical Association*, 9(70):209–219.

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

Shorrocks, A. F. (1983). Ranking income distributions. *Economica*, 50(197):3–17.

Son, H. H. (2004). A note on pro-poor growth. *Economics Letters*, 82(3):307 – 314.