Working Paper Series

Accounting for Heterogeneity in Growth Incidence in Cameroon Using Recentered Influence Function Regression

B. Essama-Nssah
Saumik Paul
Léandre Bassol

ECINEQ WP 2013 – 288

ECINEQ 2013 – 288
February 2013

www.ecineq.org

Accounting for Heterogeneity in Growth Incidence in Cameroon Using Recentered Influence Function Regression*

B. Essama-Nssah†
Independent Consultant (Former World Bank Senior Economist)

Saumik Paul
Osaka University

Léandre Bassolé
African Development Bank

Abstract

This paper frames growth incidence analysis within the logic of social impact evaluation understood as an assessment of variations in individual and social outcomes attributable to shocks and policies. It uses recentered influence function (RIF) regression to link the growth incidence curve (GIC) to household characteristics and perform counterfactual decomposition à la Oaxaca-Blinder to identify sources of variation in the distribution of consumption expenditure in Cameroon in 2001-2007. We find that the structural effect is driven mostly by the sector of employment and geography and is the main driver of the observed pattern of growth. The composition effect accounts for the lion's share of the observed variation in the social impact of growth. In particular, that effect tends to reduce poverty while the structural effect tends to increase it. This conclusion is robust with respect to the choice of poverty measures and RIF regression models.

Keywords: Cameroon, counterfactual analysis, economic growth, growth incidence curve, inequality, Oaxaca-Blinder decomposition, poverty, recentered influence function (RIF) regression, quantile regression, social evaluation.
JEL Classification: C14, C31, D31, I32, O55, R11.

---

* The authors are grateful to Francisco H. G. Ferreira, Peter J. Lambert, Rachel K. Sebudde and Abdoulaye Seck for insightful comments on earlier drafts and for encouragement. While this analysis was completed when the authors were still working for the World Bank in Washington, D.C., the views expressed herein are entirely those of the authors or the literature cited and should not be attributed to the World Bank, or to its affiliated organizations.
† **Contact details: Essama Nssah** (corresponding author), essama.nssah@gmail.com.

1. Introduction

A consensus has emerged in the development community that economic growth is not the final goal of development, but a means to the achievement of things that people ultimately care about. Sen (1989) argues that the standard definition of growth is concerned only with the amount of means of well-being available to people and is silent not only on the way those means are distributed but on what the people involved can achieve with those means given their life plans. This view was recently echoed by the Growth Commission[1] in their 2008 report. The Commissioners noted that growth is not an end in itself, but creates opportunities for the achievement of other important individual and social objectives.

If growth is a means to an end, then this end must be the yardstick by which to judge growth performance in a given country. This suggests that *growth incidence analysis* is essentially an exercise in social impact evaluation understood as an assessment of variations in individual and social outcomes induced by the process of economic growth. Variations in social outcomes are usually assessed on the basis of some social evaluation function such as a welfare function, a poverty measure or an inequality index. In general, a social evaluation function is a distributional statistic that aggregates individual outcomes into a social outcome.

Basing the evaluation of the growth process solely on changes in such a distributional statistic clearly hides more than it reveals about the heterogeneity of impacts underlying the aggregate outcome. It is well known that heterogeneity of interests and circumstances plays a central role in both policymaking and in the determination of the welfare impact of policy (Drazen 2000). Ravallion (2001) argues that, even if one were interested only in economic growth, understanding

[1] The Growth Commission is formally known as the Commission on Growth and Development. It was sponsored by the following six institutions: Australian Agency for International Development (AusAID), Dutch Ministry of Foreign Affairs, Swedish International Development Cooperation Agency (SIDA), U.K. Department for International Development (DFID), the William and Flora Hewlett Foundation, and the World Bank Group. The 21 member Commission was chaired by Michael Spence and consisted of leading policymakers mostly from developing countries, and a few academics including Robert M. Solow. The work started in April 2006 and "The Growth Report" was published in 2008.

this heterogeneity is crucial for the design and implementation of targeted interventions that might enhance the effectiveness of growth-oriented policies.

The purpose of this paper is therefore to demonstrate the use of recentered influence function (RIF) regression within the logic of Oaxaca-Blinder decomposition in order to account for heterogeneity in growth incidence with an application to the case of Cameroon. The rest of the paper is organized as follows. Section 2 presents the accounting framework based on counterfactual decomposition of the growth incidence curve (GIC) into so-called composition and structural effects. Section 3 applies the analytical framework presented in section 2 to data on Cameroon for 2001-2007 to try and understand some of the underlying factors that might have shaped the pattern of growth observed over that period of time. Section 4 contains concluding remarks.

As defined by Ravallion and Chen (2003), the GIC shows the proportional change in a living standard indicator (i.e. income or consumption expenditure) at each quantile between two points in time. While the GIC reveals some heterogeneity of impacts across quantiles, it cannot tell by itself the determining factors of this heterogeneity. Further analysis is needed in order to uncover the potential determinants of the pattern of growth as reflected by the incidence curve. To do this, we rely on counterfactual decomposition of differences across quantiles underpinning the GIC. Counterfactual decomposition of variations in a distributional statistic, such as the τ^{th} quantile (q_τ) of an outcome distribution, requires: (1) the specification of an *outcome model*; (2) an *identification* strategy; and (3) an *estimation* procedure.

The outcome model links the outcome of interest to its determining factors and thus determines the *scope* of the decomposition. In this study, we maintain that the living standard of an individual is a pay-off from her participation in the life of society. In this game-theoretic perspective, the pay-off can be viewed as a function of *participation* and *type*. Types of individuals may be characterized by their preferences, capabilities, information and beliefs (Milgrom 2004). Viewing the living standard of an individual as a function of her endowments, behavior and the circumstances that determine the returns to these endowments in any social interaction, the framework we shall discuss in section 2 will allow us to identify and estimate the contributions of endowments and returns to those endowments. Our identification strategy relies on the notion of *ceteris paribus* variation which entails the comparison of an observed outcome distribution to a counterfactual obtained by changing one factor at a time while holding all the other factors constant. We use RIF regression to estimate relevant parameters and perform both aggregate and detailed decompositions.

For the past twenty years or so, Cameroon has been battling a severe and persistent socioeconomic crisis that can be traced back, in part, to a terms-of-trade shock in the mid 1980s and the associated policy response. In October 2000, Cameroon became eligible for debt relief under the Enhanced HIPC[2] Initiative. In this context, the government adopted a Poverty Reduction Strategy (PRS) in 2003, and the country reached the Completion Point in May 2006, after three full years of implementation of the 2003 PRS. This achievement signaled the satisfaction of Cameroon's development partners with the implementation of that strategy. Between 2001 and 2007, real per capita GDP grew only by an average of 0.6 percent per year (National Statistical Office 2008). Furthermore, available data suggest that the overall incidence of poverty was 40 percent in 2007, about the same level as in 2001. The Gini index of inequality however dropped a couple of percentage points from 41 percent in 2001 to 39 percent in 2007[3].

[2] HIPC stands for Heavily Indebted Poor Countries. This initiative was launched in 1996 by the International Development Association (IDA, the World Bank's fund designed to provide concessional credits and grants to the poorest countries) and the IMF. The initiative was enhanced in 1999 to tighten its link with poverty reduction and to widen its scope and make it more efficient (in terms of speed of relief delivery). Eligibility is based on three criteria: (1) qualify only for concessional assistance from IDA, (2) debt situation remains unsustainable after full application of traditional relief mechanisms, and (3) a track record of reforms combined with the development of a Poverty Reduction Strategy (presented in a document known as Poverty Reduction Strategy Paper or PRSP). The whole process entails reaching a *Decision Point* and a *Completion Point*. Two conditions must be met by a country to reach the Decision Point: (1) satisfactory preparation of an interim PRSP, and (2) satisfactory performance under the IMF's Poverty Reduction and Growth Facility (PRGF). At this point, the country gets conditional (on continued good performance) interim relief. At the Completion Point debt relief becomes irrevocable. Reaching this point requires the following: (1) maintain macroeconomic stability under a PGRF; (2) satisfactory implementation of a full PRSP for one year; (3) implementation of structural and social reforms agreed upon at the Decision Point.

[3] See Essama-Nssah et al. (2010) for a more in depth discussion of a profile of growth, inequality and poverty in Cameroon between 1996 and 2007.

The following are the key findings emerging from this study. The composition (or endowment) effect accounts for the level of the growth incidence curve, and for the lion's share of the observed variation in social outcomes. The structural effect is the main factor explaining the shape of the observed pattern of growth. Household demographics are the key factor driving the composition effect while the sector of employment and geography account for most of the structural effect. Finally, the urban bias that characterizes the pattern of growth has increased over time.

2. Accounting Framework

Within the logic of social evaluation the GIC can be interpreted as a social impact indicator. This section reviews the structure and the normative underpinnings of the growth incidence curve along with the decomposition of the variation in a distributional statistic into the *composition* and *structural* effects. Finally, it explains the use of RIF regression for growth incidence analysis.

2.1. The Growth Incidence Curve

Structure

Let y be the outcome variable of interest, say an indicator of economic welfare such as income or consumption expenditure. Following a growth episode between t=0 and t=1, the incidence of growth on the income of individual i can be written as follows.

$$g(y_i) = \left(\frac{y_{1i}}{y_{0i}} - 1\right) \tag{2.1}$$

where $g(y_i)$ is the rate of change of individual i's income following the growth spell.

We cannot identify and hence compute this individual level effect unless we have panel data spanning the growth episode. This is a manifestation of the missing data problem that arises in the context of program impact evaluation. In the absence of panel data, we rely on the assumption of *anonymity* (or *symmetry*) to identify and estimate the above effect on the basis of a local impact indicator that compares the initial and posterior distributions of income across quantiles[4]. Anonymity implies that, when comparing distributions of living standards, the position of a particular individual in one distribution is irrelevant (Carneiro, Hansen and Heckman 2002). In other words, the identity of the income recipients is irrelevant for this comparison[5]. Thus a permutation of the incomes of any two individuals in any of the two distributions being compared does not affect the comparison.

Let's assume that income is continuously distributed over the population of interest. We denote by $F_t(y)$ the cumulative distribution function (CDF) of income showing the proportion, τ, of the population with income less than y at time t. The income level at the τ^{th} quantile[6] is given by the inverse of the CDF: $y_t(\tau) = F_t^{-1}(\tau)$. The growth rate of income at the τ^{th} quantile between t=0 and t=1 is equal to the following expression.

$$g(\tau) = \left[\frac{y_1(\tau)}{y_0(\tau)} - 1\right] \qquad (2.2)$$

The GIC as defined by Ravallion and Chen (2003) is obtained by letting τ vary from zero to one and plotting the corresponding values of $g(\tau)$. The quantiles involved in the computation of equation (2.2) are based on the ranking of individuals according to their baseline income. This is a consequence of the veil of ignorance (anonymity) shrouding the comparison of the two distributions. To make this point abundantly clear, following Essama-Nssah and Lambert (2009), if we let $F_0(y)$ and $F_1(v)$ be

[4] In the context of program impact evaluation (or treatment effect analysis) the missing data problem stems from the impossibility of observing simultaneously the same individual in two different states of nature (one with and the other without treatment). Observing the same individual before and after treatment does not usually solve this identification problem because of potential confounding factors. However, before and after comparisons identify impact in the context of growth incidence analysis because, in addition to anonymity, we assume that changes in all factors affecting individual incomes show up through their effects on the overall growth process.

[5] This assumption plays the same role as the one played by the assumption of *rank preservation* for the identification of quantile treatment effects in the context of treatment effect analysis. Rank preservation (or *rank invariance*) across two alternative states of the world means that the outcome at the τ^{th} quantile of the outcome distribution in one state has its counterpart at the same quantile of the outcome distribution of the other state. When rank preservation fails, the approach identifies and estimates the difference between the quantiles and not the quantiles of the difference in outcomes (Bitler et al. 2006).

[6] Let f(·) stand for the density function characterizing the distribution of y, then $\tau = \int_0^y f(x)dx$.

respectively the distribution functions before and after growth takes place then equation (2.2) becomes[7].

$$g(\tau) = \left(\frac{F_1^{-1}(F_0(y))}{y} - 1\right), \ \tau = F_0(y) \qquad (2.3)$$

Both expressions (2.2) and (2.3) clearly show that the GIC compares the posterior distribution with the baseline distribution of y over their entire domain of definition. In other words, the curve shows the growth rate of each quantile of the distribution. The curve can thus be considered as an indicator of the *pattern of growth*. Taking the distribution of income at a given date as an indicator of the corresponding social state, we can also view the GIC as an indicator of the social impact of growth. It provides a factual (as opposed to normative) description of change in inequality associated with the growth process. If the GIC is a decreasing function for all τ in its domain of definition, then all inequality measures that respect the Pigou-Dalton principle of transfers will indicate a fall in inequality over time. If instead, the GIC is an increasing function of τ, then the same measures will register an increase in inequality (Ravallion and Chen 2003). When inequality does not change the GIC will show the same growth rate for all τ.

Figure 2.1 presents the growth incidence curve for Cameroon for the period 2001-2007. This curve shows how the distribution of per equivalent adult expenditure[8] changes at each quantile between 2001 and 2007. Presumably this is an outcome of the underlying Poverty Reduction Strategy the implementation of which started in 2003. The curve reveals some *heterogeneity* in the impact of growth on the living standards. People located at the bottom of the distribution up to the 10th percentile have experienced an income growth greater than average and so have most of the people above the median, except at the very top of the

[7] Similarly, Bourguignon (2010) defines the GIC associated with a growth path from distribution F(·) to distribution H(·) as: $g_{HF}(\tau) = \frac{y_H(\tau)}{y_F(\tau)} - 1$.

[8] The underlying scale assigns weights to individual members of the household according to their age and gender. However there is no gender differential for children up to the age of 10. Thus children who are at most 1 year old get a weight of 0.255. Those with age between 1 and 3 years get assigned a weight of 0.45. Between the age of 4 and 6, the weight is 0.62 while it is 0.69 for the 7-10 age group. Starting from age 11, males get assigned the following weights: 0.86 between 11 and 14, 1.03 between 15 and 18, 1 between 19 and 50 and 0.79 above 50. All females between 11 and 50 get a weight of 0.76 and those above 50 get a weight of 0.66.

distribution. Between the 10th and about the 30th percentiles, incomes grew at a rate below average. Finally the segment of the population located between the 30th and the 50th percentiles experienced an income growth rate equal to the growth rate of the average income. On the basis of changes in the slope of this GIC, we conclude that inequality has fallen both at the very bottom and the very top of the distribution, and has risen somewhat in the middle.

Figure 2.1.Growth Incidence Curve for Cameroon, 2001-2007

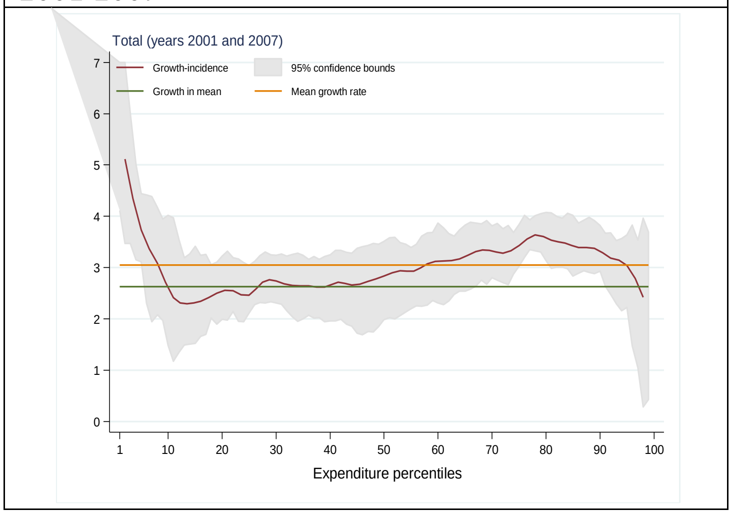

One can split the GIC into one component showing the growth rate of average income, γ, and the other showing for each point on the curve, the deviation from the overall growth rate. Formally, we have: $g(\tau) = \gamma + [g(\tau) - \gamma]$. The first component is essentially a scale factor. This is the rate of growth that would be experienced at every quantile if growth were distribution neutral. The second component shows what the incidence would be if the growth process changed only relative inequality and not the mean income. In this case, the process would be purely redistributive. We refer to the first component as the *size effect* and the second as the *redistribution effect* of growth. This decomposition makes it clear that the pattern of growth is determined mainly by the second component (the redistribution effect). Ravallion and Chen (2003) provide an alternative expression for this decomposition as: $g(\tau) = \gamma + \Delta lnL'(\tau)$, where the redistribution effect is written in terms of changes in the slope of the relevant Lorenz curve[9]. This expression suggests that the GIC will be greater than the overall growth rate if and only if the slope of the Lorenz curve increases over time.

Figure 2.2 Food and Nonfood Components of Growth Incidence in Cameroon
(2001-2007)

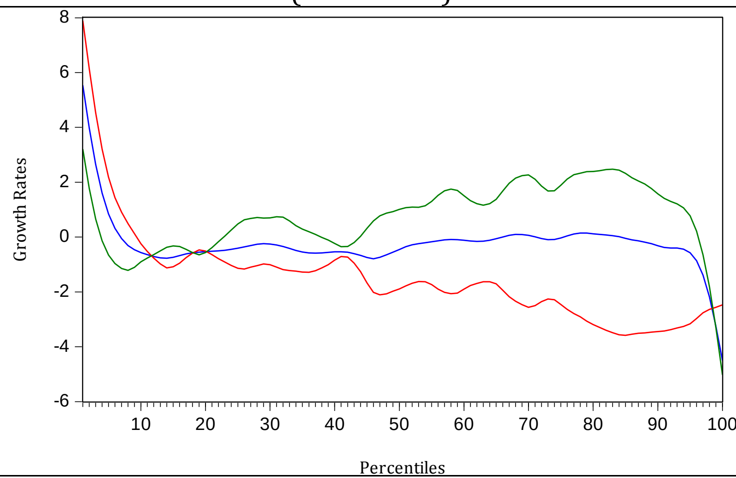

Essama-Nssah and Lambert (2009), show that the GIC is also decomposable across components of a living standard indicator. To see this, suppose that income y has two components, such that: $y = y_1 + y_2$. Then the overall GIC can be written as a weighted average of the growth rates of the components. The weights are equal to income shares. The overall growth rate at the τ[th] quantile become:

$$g(\tau) = \frac{y_1}{y}\frac{\Delta y_1}{y_1} + \frac{y_2}{y}\frac{\Delta y_2}{y_2} = s_1 g_1(\tau) + (1 - s_1) g_2(\tau) \qquad (2.4)$$

where $s_1 = \frac{y_1}{y}$. If there are more than two income sources, say m, the decomposition generalizes to: $g(\tau) = \sum_{j=1}^{m} s_j g_j(\tau)$ where $s_j = \frac{y_j}{y}$, $g_j(\tau) = \frac{\Delta y_j}{y}$ and $\sum_{j=1}^{m} s_j = 1$.

Figure 2.2 shows the contribution of food and nonfood components of household expenditure to the pattern of growth depicted in figure 2.1. It can be seen from this graph that both components of expenditure reduce inequality at the

[9] Let L(τ) be the ordinate of the Lorenz curve at the τ[th] quantile. It is well known that the first-order derivative of the Lorenz curve at this point is equal to the level of income at that quantile divided by the overall mean of the distribution: $L'(\tau) = \frac{y(\tau)}{\mu}$. This fact implies that $y(\tau) = \mu L'(\tau)$. Computing the growth rate of income at this quantile as the change in the logarithm of the corresponding income level over time yields the desired expression.

bottom of the distribution, up to the 20th percentile. The two curves begin to diverge beyond this point. While changes in food expenditure continue to have a dampening effect of overall inequality, the nonfood component tends to increase inequality up to the neighborhood of the 90th percentile. The shape of the overall GIC suggests the impact of the nonfood component dominates that of the food component beyond the 20th percentile.

Normative Considerations

As noted earlier, economic growth is not an end, but a means to achieving what individuals and society care about. We suppose that the extent to which an individual cares about her income level y can be stated by u(y), the level of utility associated with income y. This utility function represents the welfare of an individual with income level y. Given a profile of individual outcomes, society's concerns can be represented by a social welfare function of the form, W(u(y)).

Concerns about inclusive growth can be modeled by the class of second-order social welfare functions, denoted by Ω_2. This class of social welfare functions agree with both individual preference for prosperity (more is preferred to less) and social preference for equality. It is included in the class of first-order social welfare functions, Ω_1, which approves of individual preference for prosperity but is silent on social preference for equity. Formally we note: $\Omega_1 = \{W(u(y)) | u'(y) \geq 0\}$. According to Duclos and Araar (2006), the second-order class, Ω_2, includes all members of Ω_1 that also approve of *mean-preserving equalizing transfers*[10]. Such transfers are consistent with the Pigou-Dalton principle of transfers according to which a transfer of income from a richer to a poorer person increases social welfare, *ceteris paribus*. This idea can be implemented by making the utility of income a strictly concave function. The second-order class of social welfare functions can therefore be characterized as follows: $\Omega_2 = \{W(u) | W \epsilon \Omega_1, u''(y) \leq 0\}$. Thus second-

[10] To see clearly what is involved here, consider two individuals j and k with income levels $y_j \leq y_k$ so that k is richer than j. Now, *ceteris paribus*, take a small amount of income from k, say δ, and give it to j so that k is still richer than j afterwards. The transfer is mean-preserving because total income (hence the mean) does not change. It is also equalizing because it reduces inequality between these two individuals. See Duclos and Araar (2006) or Moulin (2003) for details. Social welfare functions that respect the Pigou-Dalton principle are said to be *distribution-sensitive*.

order social judgments entail further restrictions on members of the first-order class of social welfare functions[11].

The growth incidence curve can reveal dominance relations among the underlying income distributions. When this curve is greater or equal to zero for all τ (i.e. $g(\tau) \geq 0, \forall \tau$), the posterior distribution function of income lies nowhere above and at least somewhere below that of the initial distribution. This relation is known as *first-order stochastic dominance* expressed by the relation: $F_1(y) \leq F_0(y), \forall y$. This is equivalent to saying that all members of the first-order class of social welfare functions will declare the move from the initial to the posterior distribution a social improvement.

First order stochastic dominance fails when the GIC switches sign over its domain of definition. The two underlying distributions can still be compared but now on the basis of *second order stochastic dominance* using the associated deficit functions. The deficit function is the integral of a CDF over its domain. If the deficit function for the posterior distribution lies nowhere above and somewhere below that associated with the initial distribution, then the former dominates the latter to the second order. First order stochastic dominance implies second order stochastic dominance, but not the other way around.

The extent to which growth is inclusive can also be assessed on the basis of poverty implications of the pattern of growth. Grimm (2007) advocates dropping the anonymity postulate in the context of pro-poor growth analysis. He argues that it is of policy interest to determine whether a given growth experience is helping the initially poor get out of poverty and maybe pushing some originally non-poor back into poverty. To be able to account for such mobility in and out of poverty, he proposes a modified growth incidence curve which he calls "*Individual Growth Incidence Curve*". This curve is computed on the basis of panel data. Individuals are ordered according to their position in the initial distribution. Holding membership to initial quantiles fixed, the outcomes for individuals in each such quantile are

[11] See Duclos and Araar (2006) for a discussion of higher-order welfare dominance.

identified[12] in the final distribution and used to compute quantile-specific means and growth rates. The author applies this approach to panel data for Indonesia and Peru and finds a stronger effect of economic growth on the initially poor than suggested by cross-section evidence.

Bourguignon (2010) also argues that growth incidence analysis based on standard growth incidence curves ignores income mobility and is tantamount to assuming that only post-growth income matters in social evaluation. Given anonymity, a standard GIC compares the income of individuals who were not necessarily in the same position in the initial distribution. Re-ranking may have occurred in the movement from the initial to the posterior distribution. This author therefore recommends the use of *non-anonymous* growth incidence curves to base the evaluation on both initial and terminal incomes.

A non-anonymous GIC plots the rates of growth on income against quantiles for the initial distribution of income. This curve is derived as follows. Let y stand for initial period income with a distribution function F(·) and support $(0, m_y)$. Let z stand for end period income with distribution H(.). Fix the ranking of income units according to the initial distribution. A non-anonymous GIC links each initial quantile, y(τ), to a corresponding one in the posterior distribution. This curve is formally defined by the following expression.

$$g_{na}(\tau) = \frac{\int_0^{m_y} z dH(z|y(\tau))}{y(\tau)} - 1 \qquad (2.5)$$

Where H(z|y) stands for the distribution of terminal income conditional on initial income. Thus a given point of the non-anonymous GIC measures the mean income growth of all units located at the τ^{th} quantile of the initial distribution. Just as in the case of the individual growth incidence curve of Grimm (2007), the computation of non-anonymous GICs requires panel data.

2.2. The Composition and Structural Effects

[12] This identification is made possible by the panel nature of the data. With cross-sectional data, one has no choice but to resort to the anonymity postulate or to microsimulation techniques.

It is useful to identify the factors shaping a pattern of growth. Counterfactual decomposition of distributional change provides a way of doing so[13]. The analysis can be framed within the logic of Oaxaca-Blinder decomposition[14] based on an outcome model that considers both individual and social outcomes where the social outcome is represented by a distributional statistic, say θ, viewed as a functional[15] of the distribution of individual outcomes (F_y), and we take an individual outcome to be a function of her endowments, behavior and the circumstances that determine the returns to her endowments in any socioeconomic transaction. In this section, we consider both aggregate and detailed decompositions, noting the implications for the class of additively separable social evaluation functions.

The Outcome Model

We are interested in decomposing a change in some distributional statistic, say θ, from the base period t=0 to the end period t=1. Let $F_{y_0|t=0}$ stand for the outcome distribution observed in the initial period and $F_{y_1|t=1}$ the distribution for the final period. The distributional change from state 0 to 1 can be characterized by the variation in $\theta(F_y)$ as follows:

$$\Delta_O^{\theta} = \theta\left(F_{y_1|t=1}\right) - \theta\left(F_{y_0|t=0}\right) \tag{2.6}$$

We seek to decompose this overall difference on the basis of the relationship between the outcome variable and individual or household characteristics. The following equation represents a general expression of that relationship.

---

[13] This presentation draws extensively on Firpo, Fortin and Lemieux (2009) and Fortin, Lemieux and Firpo (2011). These authors pioneered the RIF regression approach to counterfactual decomposition.

[14] The classic Oaxaca-Blinder method seeks to decompose the overall difference in *unconditional mean outcome* between two groups or time periods into two components: (1) the composition or endowment effect due only to changes in the distribution of observable characteristics, and (2) the structural or price effect due to changes in returns to those characteristics. This method uses a linear regression model to link the outcome variable to individual characteristics and assumes that the conditional expectation of the error term given the observables is equal to zero. This assumption, combined with the law of iterated expectations implies that the coefficient associated with a covariate in the regression model measures both the effect of the covariate on the conditional mean outcome and the effect of increasing the mean of the covariate on the unconditional mean outcome. The latter motivates the classic Oaxaca-Blinder approach which has now been extended to general distributional statistics (besides the mean) and to entire outcome distributions.

[15] Roughly speaking, a functional is a function of a function. In this particular context, it is a rule that maps every outcome distribution in its domain into a real number (Wilcox 2005).

$$y_t = \varphi_t(x_t, \varepsilon_t),\ t = 0, 1. \qquad (2.7)$$

Identification

Equation (2.7) suggests that conditional on the observable characteristics, x, the outcome distribution depends only on the function $\varphi_t(\cdot)$ and the distribution of unobservable characteristics, ε. The fact that the distributional statistic of interest depends on the distribution of y suggests there are four potential terms for the decomposition of (2.6). In other words, differences in outcome distributions between the two periods may be due to: (i) differences in returns to observable characteristics given the functions defining the outcome structure, (ii) differences in returns to unobservable characteristics also defined by the structural functions, (iii) differences in the distribution of observable characteristics, and (iv) differences in the distribution of unobservable characteristics.

Under the general outcome model presented in equation (2.7), it is impossible to distinguish the contribution of returns to observables from that of returns to unobservables. These two terms are therefore lumped in a single term, the structural effect noted Δ_S^θ. Let Δ_X^θ stand for the endowment (or composition) effect and $\Delta_\varepsilon^\theta$ for the effect associated with differences in the distribution of unobservables. These terms are usually indentified through restrictions that ensure that each term emerges from a *ceteris paribus* variation of the relevant factors. Let $y_{0|t=1}$ be the outcomes that would have prevailed in period 1 if individual characteristics in that period had been rewarded according to $\varphi_0(\cdot)$. Let $F_{y_0|t=1}$ stand for the corresponding distribution and $\theta\left(F_{y_0|t=1}\right)$ for the corresponding value of the statistic of interest. We identify the composition effect by:$\Delta_X^\theta = \left[\theta\left(F_{y_0|t=1}\right) - \theta\left(F_{y_0|t=0}\right)\right]$ since the *ceteris paribus* condition must hold for both the conditional distribution of unobservable (given the observables) and the outcome structure. The assumption of ignorability commonly used in observational studies in the context of impact evaluation requires that the conditional distribution of unobservable factors be the same in both states of the world, hence $\Delta_\varepsilon^\theta = 0$. If we further assume that the outcome structure, φ(·), remains stable[16] as we adjust the distribution of observables to obtain the relevant counterfactual outcome, then the structural effect is due solely to differences in the functions defining the outcomes[17]. This effect is identified by: $\Delta_S^\theta = \left[\theta\left(F_{y_1|t=1}\right) - \theta\left(F_{y_0|t=1}\right)\right]$.

To summarize, given the outcome model represented by equation (2.7), assuming mutually exclusive groups, common support, simple counterfactual treatment and ignorability, we can decompose the distributional difference in equation (2.6) by adding to and subtracting from it the counterfactual outcome $\theta\left(F_{y_0|t=1}\right)$:

$$\Delta_O^\theta = \left[\theta\left(F_{y_0|t=1}\right) - \theta\left(F_{y_0|t=0}\right)\right] + \left[\theta\left(F_{y_1|t=1}\right) - \theta\left(F_{y_0|t=1}\right)\right] \quad (2.8)$$

In which the first term on the right hand side is the *composition effect* and the second is the *structural effect*.

Estimation

DiNardo, Fortin and Lemieux (1996) show that the counterfactual distribution, $F_{y_0|t=1}$, can be estimated by properly reweighting the distribution of covariates in period 0. One can express the resulting counterfactual distribution as follows[18].

$$F_{y_0|t=1}(y) = \int F_{y_0|x_0}(y|x)\mathrm{w}(x)dF_{x_0}(x) \quad (2.9)$$

---

[16] This assumption is some time referred to as simple treatment assumption (or no general equilibrium effects).

[17] To see this, note that $y_{1|t=1} = \varphi_1(x_1, \varepsilon_1)$ and $y_{0|t=1} = \varphi_0(x_1, \varepsilon_1)$.

[18] The process of reweighting adjusts the distribution of the covariates x in period t=0 so that it becomes similar to that in period t=1. For this adjustment to help us identify the terms of the decomposition, it must be a *ceteris paribus* adjustment. Since $y_0 = \varphi_0(x, \varepsilon)$, the ceteris paribus condition would be violated if changing the distribution of x also changed either the function $\varphi_0(\cdot)$ or the conditional distribution of ε given x. This would confound the impact of the adjustment and the decomposition would be meaningless. Changes in the structural function are ruled out by the simple treatment assumption (no general equilibrium effects) while changes in the conditional distribution of ε are ruled out by the ignorability assumption. Under this circumstances, we expect the conditional distribution of y_0 given x to be invariant with respect to adjustments in the distribution of the observable factors x.

where the reweighting factor is: $\mathrm{w}(x)=\frac{dF_{x_1}(x)}{dF_{x_0}(x)}=\frac{P(t=1|x)}{1-P(t=1|x)}\cdot\frac{1-\pi}{\pi}$. These weights are proportional to the conditional odds of being observed in state 1. The proportionality factor depends on π which is the proportion of cases observed in state 1. One can easily compute the reweighting factor on the basis of a probability model such as logit or probit. Furthermore, if one is interested only in the aggregate decomposition of the variation in a distributional statistic, then all that is needed are an estimate of the relevant counterfactual distribution and the corresponding value of the statistic in question. The decomposition presented in equation (2.8) is based on nonparametric identification and can be estimated by the Inverse Probability Weighing method implied by equation (2.9).

A decomposition approach provides a detailed decomposition when it allows one to apportion the composition effect or the structural effect into components attributable to each explanatory variable. The contribution of each explanatory variable to the composition effect is analogous to what Rothe (2010) calls a "*partial composition effect*"[19]. As discussed earlier, this is easily accomplished in the context of the classic Oaxaca-Blinder decomposition because of the two underlying assumptions of linearity and zero conditional mean for the unobservable factors.

Recentered influence function regression (henceforth, RIF regression) offers a simple way of establishing a direct link between a distributional statistic and individual (or household) characteristics. This link offers an opportunity to perform both aggregate and detailed decompositions for any such statistic for which one can compute an *influence function*. In essence, the influence function of a functional θ(F) is its first-order directional derivative (Hampel 1974). To see this, let G(b) be a mixture of two distributions F and H such that an observation is randomly sampled from F with probability b or from H with probability (1-b). In other words, $G(b)=bH+(1-b)F$. We are interested in how the functional θ changes as G gets closer

[19] This is the effect of a counterfactual change in the marginal distribution of a single covariate on the unconditional distribution of an outcome variable, *ceteris paribus*. Rothe (2010) interprets the *ceteris paribus* condition in terms of rank invariance. In other words, the counterfactual change in the marginal distribution of the relevant covariate is constructed in such a way that the joint distribution of ranks is unaffected.

and closer to F. The effect of this change is revealed by the directional derivative of θ at F in the direction of H. Formally, we write:

$$\nabla\theta_{G\to F} = \lim_{b\to 0}\frac{\theta(G(b))-\theta(F)}{b} = \frac{\partial}{\partial b}\theta(bH+(1-b)F)|_{b=0} \tag{2.10}$$

Consider the particular case where $H = \Delta_y$,a distribution function for a probability measure that assigns mass 1 to y, in the domain of F. In other words, $\Delta_y(v)$ is equal to 1 if $y \leq v$, otherwise[20] it is equal to 0. In this case, $G = b\Delta_y + (1-b)F$ and the influence function of the functional θ(F) can be written as:

$$IF(y;\theta,F) = \nabla\theta_{F\to\Delta_y} \tag{2.11}$$

This is a measure of the relative effect of a small perturbation in F on θ(F). In that sense, it is a measure of robustness[21]. The influence function defined in (2.11) measures the effect that a single observation has on a functional. Since:

$$\frac{\partial}{\partial b}\left[\int v\big(bh(v)+(1-b)f(v)\big)dv\right]_{b=0} = \int vh(v)dv - \int vf(v)dv \tag{2.12}$$

In other words, the directional derivative of the mean of F at F in the direction of H is equal to: $\nabla\mu_{\mathrm{F}\to\mathrm{H}} = \mu_{\mathrm{H}} - \mu_{\mathrm{F}}$. In particular, the influence function of the mean of F is equal to[22]:

$$IF(y;\mu,F) = \nabla\mu_{F\to\Delta_y} = y - \mu_F \tag{2.13}$$

[20] This can be expressed with an indicator function as follows: $\Delta_y(v) = I(y \leq v)$. Recall that an indicator function is equal to one when its argument is true and zero otherwise. In particular $\int_0^\infty \Delta_y(v)f(y)dy = \int_0^\infty I(y\leq v)f(y)dy = \int_0^v f(y)dy = F(v)$.

[21] Wilcox (2005) explains that continuity alone confers only qualitative robustness to the statistic under consideration. A continuous function is relatively unaffected by small shifts in its argument. Similarly, differentiability is related to infinitesimal robustness in the sense that, if a function is differentiable and its derivative is bounded, then small variations in the argument will not result in large changes in the function. Thus a search for robust statistics can focus on functionals with bounded derivatives.

[22] Influence functions can also be computed for many other distributional statistics. Essama-Nssah and Lambert (2012) show how to derive the influence function of a functional from the associated directional derivative. In addition, they present a collection of influence functions for social evaluation functions commonly used in assessing the social impact of public policy. Their catalog includes, among others, influence functions and recentered influence functions for the mean, the τ^{th} quantile, the Gini coefficient, the Atkinson index of inequality, the class of additively separable poverty measures, the growth incidence curve ordinate, the Lorenz curve and generalized Lorenz curve ordinates, the TIP curve ordinate and some measures of pro-poorness associated with the Foster, Greer and Thorbecke (1984) family of poverty measures.

The expected value of the influence function of a distributional statistic is in fact equal to zero in all cases in which the frequencies and the range of the y-values are bounded.

Firpo, Fortin and Lemieux (2009) define the recentered or rescaled influence function (RIF) as the leading two terms of a von Mises (1947) linear approximation of the associated functional[23], namely as the functional itself plus the corresponding influence function. Letting IF(y; θ) stand for the influence function of $\theta(F_y)$, $RIF(y;\theta) = \theta(F_y) + IF(y;\theta)$.

The fact that the expected value of the influence function is equal to zero implies that the expected value of the RIF is equal to the corresponding distributional statistic: $\theta(F_y) = E[RIF(y;\ \theta)]$. By the law of iterated expectations the distributional statistic of interest can be written as the conditional expectation of the rescaled influence function (given the observable covariates, x). This is the RIF regression that, for $\theta(F_y)$, can be expressed as: $E[RIF(y;\theta)|x]$. The distributional statistic $\theta(F_y)$ can therefore be expressed in terms of this conditional expectation as follows.

$$\theta(F_y) = \int E[RIF(y;\ \theta)|x]dF(x) \tag{2.14}$$

To assess the impact of covariates on $\theta(F_y)$, one needs to integrate over the conditional expectation $E[RIF(y;\theta)|x]$, which can be done using regression methods. If we model this conditional expectation as a linear function of observable covariates, $E[RIF(y;\theta)|x] = x\beta$, then we may apply OLS to the following equation.

$$RIF(y;\theta) = x\beta + \varepsilon \tag{2.15}$$

in order to estimate β.

The expected value of the linear approximation of the RIF regression is equal to the expected value of the true conditional expectation because the expected value

[23] This is analogous to the approximation of a differentiable function at a point by a Taylor's polynomial.

of the approximation error is zero. This fact makes the extension of the standard Oaxaca-Blinder decomposition to RIF regression both simple and meaningful[24].

Applying the standard Oaxaca-Blinder approach to equation (2.15) we find that the composition effect can be written as follows.

$$\Delta_X^{\theta} = [E(x|t=1) - E(x|t=0)] \cdot \beta_0 \tag{2.16}$$

The corresponding structural effect is

$$\Delta_S^{\theta} = E(x|t=1) \cdot (\beta_1 - \beta_0) \tag{2.17}$$

where β_0 and β_1 are returns to observable endowments in periods 0 and 1 respectively.

2.3. Use of RIF Regression for Growth Incidence Analysis

One can exploit the linearity of the RIF regression model presented in equation (2.15) to perform both aggregate and detailed decompositions of differences across quantiles, and thus decompose growth incidence into the composition and structural effects and identify the contribution of covariates to these effects. One would need recentered influence functions for the quantiles under consideration. The rescaled influence function of the τ^{th} quantile of the distribution of y is the following[25]:

$$RIF(y;\ q_\tau) = q_\tau + IF(y;\ q_\tau) = q_\tau + \frac{[\tau - I(y \leq q_\tau)]}{f_y(q_\tau)} \tag{2.18}$$

Where $I(\cdot)$ is an indicator function for whether the outcome variable y is less than or equal to the τ^{th} quantile and $f_y(q_\tau)$ is the density function of y evaluated at the τ^{th} quantile.

One can use equation (2.18) repeatedly to decompose, for instance, the first 99 quantiles (percentiles) of the outcome distribution of interest. This means that we can decompose the growth incidence curve (GIC) into a component associated

[24] The influence function for the mean presented in equation (2.13) implies that the corresponding recentered influence function is RIF(y;μ, F)=y. Hence, the ordinary linear regression of y on a set of covariates x is indeed a RIF regression.

[25] To see where this expression comes from, let q_τ be the τ^{th} quantile of F. Also, let $q_\tau(b)$ stand for the τ^{th} quantile of the mixed distribution G(b) so that $G(b) = bH(q_\tau(b)) + (1-b)F(q_\tau(b)) = \tau$, and $q_\tau(0) = q_\tau$. The first-order derivative of G with respect to b, evaluated at b=0, yields the following expression: $H(q_\tau) - F(q_\tau) + f(q_\tau)q'_\tau(0) = 0$. Hence the directional derivative of this quantile in the direction of H is equal to: $q'_\tau(0) = \frac{F(q_\tau) - H(q_\tau)}{f(q_\tau)} = \frac{\tau - H(q_\tau)}{f(q_\tau)}$. Setting $H(q_\tau) = \Delta_y(q_\tau) = I(y \leq q_\tau)$ implies that the influence function of the τ^{th} quantile of F is equal to: $IF(y; q_\tau, F) = \frac{\tau - I(y \leq q_\tau)}{f(q_\tau)}$. See Essama-Nssah and Lambert (2012) for more details.

with the composition (or endowment) effect and a second one related to the structural (or the price) effect. Formally,

$$g(y) = g_X(y) + g_S(y) \qquad (2.19)$$

Each of these components can further be decomposed in terms of the contributions of the relevant covariates.

If both first- and second-order dominance fail to hold between the initial and posterior outcome distributions, it is common to assess distributional changes on the basis of the value judgments underlying specific social evaluation functions. Here we consider the decomposition of variations in social welfare functions and poverty outcomes.

Social Welfare

To see what is involved, we consider *utilitarian social welfare functions* that are members of the second order class, Ω_2, for which individuals and society care about both *individual prosperity* and *social progress* (understood as a reduction in inequality). These values are consistent with the idea of inclusive growth.

Utilitarianism is indeed the most widely used method of passing social judgments in which social welfare is measured by the sum of individual utilities. We therefore define social welfare W simply as average utility in society (Lambert 2001):

$$W = \int_0^{m_y} u(y)f(y)dy = \int_0^{m_y} u(y)dF(y) \qquad (2.20)$$

For this class of social welfare functions, the elasticity of marginal utility of income reveals the degree of inequality aversion, defined as: $\eta_u = -\frac{yu''(y)}{u'(y)}$. The change in social welfare induced by the pattern of growth defined by an incidence curve $g(y)$ is given by the following expression (Essama-Nssah and Lambert 2009):

$$\Delta W = \int_0^{m_y} yu'(y)g(y)dF(y) \qquad (2.21)$$

where m_y stands for the maximum level of living standard indicated by y. Thus, such functions inherit the decomposability of the growth incidence curve, and the corresponding variations in social outcomes can be decomposed along the same dimensions as those described above for the GIC. First-order dominance of the posterior distribution over the initial one, for which $g\big(y(\tau)\big) \geq 0, \forall\tau$, means that social welfare functions defined by (2.20) will register an improvement for every increasing utility function u(y). Recall that marginal utility of income is positive since individuals prefer more to less.

Applying (2.19) to (2.21), the overall change in social welfare can now be expressed as $\Delta W = \Delta_X W + \Delta_S W$, where the composition effect is:

$$\Delta_X W = \int_0^{m_y} yu'(y) g_X(y) dF(y) \tag{2.22}$$

and the structural effect is:

$$\Delta_S W = \int_0^{m_y} yu'(y) g_S(y) dF(y) \tag{2.23}$$

We can now compute the contribution of different subsets of covariates to the composition and structural effects. Let $g_{Xk}(y)$ be the contribution of covariate k to the composition component of the GIC, and $g_{Sk}(y)$ its contribution to the structural component. The corresponding elements for the change in social welfare are

$$\Delta_{Xk} W = \int_0^{m_y} yu'(y) g_{Xk}(y) dF(y) \tag{2.24}$$

which is the contribution of characteristic k to the composition effect, and

$$\Delta_{Sk} W = \int_0^{m_y} yu'(y) g_{Sk}(y) dF(y) \tag{2.25}$$

for the structural effect.

For empirical implementation, one can use the *Atkinson social welfare function* for which: $u(y) = \frac{y^{1-\varepsilon}}{1-\varepsilon}, \varepsilon \neq 1$ which belongs to Ω_2 if $\varepsilon \geq 0$ which a measures the degree of inequality aversion (Atkinson 1970). When $\varepsilon=1$, the utility function takes the form $u(y) = ln(y)$.

When $\varepsilon=0$, the social welfare function is insensitive to inequality: social evaluation is indifferent between increasing the income of a poor person by a certain amount and increasing a richer person's income by the same amount, being based on prosperity alone: $\Delta W = \Delta\mu = \int_0^{m_y} yg(y) dF(y)$. The rate of growth, γ, can be written as a weighted sum of points along the growth incidence curve[26]:

$$\gamma = \frac{\Delta\mu}{\mu} = \int_0^{m_y} \frac{y}{\mu} g(y) dF(y) \tag{2.26}$$

[26] Ferreira (2010) derives the same expression through direct differentiation.

This expression confirms that the rate of growth of per capita income is a weighted average of points on the GIC where the weights are given by the slope of the Lorenz curve (a local indicator of inequality). It is clear from expression (2.26) that the decompositions discussed above apply also to this case.

Figure 2.3.Weighing System for the S-Gini Welfare Function

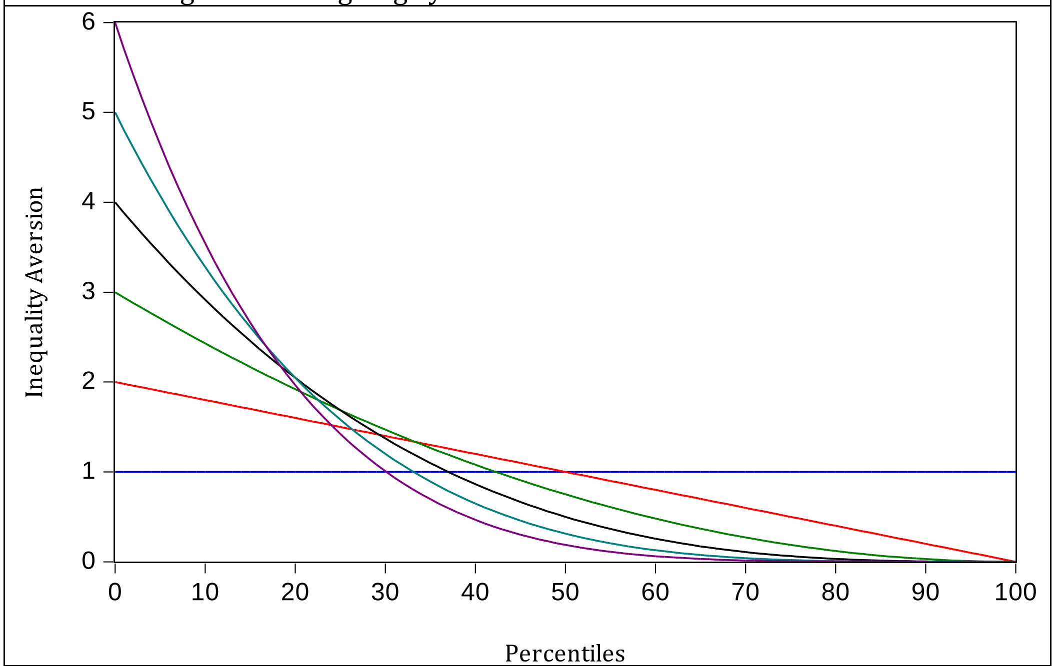

If one believes that the social value of an individual welfare should depend on the relative position of that individual in the overall distribution of income, one can use the following member of Ω_2:

$$W = \int_0^{m_y} y\omega\big(F(y)\big)dF(y); \;\; \omega(F(y)) = \nu[1 - F(y)]^{\nu-1} \qquad (2.27)$$

The weighing function, ω(·) is the one underlying the extended Gini coefficient defined by Yitzhaki (1983) and is dependent on rank since F(y) is the proportion of individuals with income less than or equal to y. Duclos and Araar (2006) call this welfare function *S-Gini social welfare function* (S-Gini is short for single parameter Gini). These authors interpret the social welfare function as the expected income of the poorest individual in a group of ν randomly selected individuals. As in the case of the Atkinson social welfare function, ν is a measure of aversion to inequality. The higher its value, the more weight is assigned to individuals at the bottom of the distribution relative to elsewhere.

Figure 2.3 show level curves characterizing the weighing function underlying the S-Gini social welfare function. When ν =1 everybody receives equal consideration regardless of rank. This is a benchmark case in the sense that, this horizontal level defines for each value of ν a cut-off rank where weights switch from greater than one to less than one as we move from the lower end to the upper end of the distribution (Essama-Nssah 2002). Over this trajectory, the weights fall progressively from the highest value for the poorest individual to zero for the richest. When the aversion parameter is equal to two, the weighing system is consistent with that of the ordinary Gini coefficient. In this case, as figure 2.3 reveals, the cut-off point is located at the median. The weights fall at a constant rate from 2 to 0 in this case.

Under the assumption of rank invariance (anonymity), the change in social welfare induced by a growth pattern $g(y)$ for this evaluation function is

$$\Delta W = \int_0^{m_y} \nu[1 - F(y)]^{\nu-1} g(y) dF(y) \quad (2.28)$$

Essama-Nssah (2005) interprets this social impact indicator as the equally distributed equivalent growth rate. This is the growth rate that will be socially equivalent to the observed one for the given choice of the degree of inequality aversion. Because this indicator is a weighted average of points on the growth incidence curve, it too is amenable to the decompositions presented above.

Poverty

For the class of additively separable poverty measures of the form

$$P(F; z) = \int_0^{m_y} I(y \leq z) \psi(y|z) dF(y) \quad (2.29)$$

where z is the poverty line and I(y≤z) is an indicator function which is equal to zero when the welfare indicator is greater than the poverty line, a change in poverty over time can also be written as a weighted sum of points on the growth incidence curve (Essama-Nssah and Lambert 2009, Ferreira 2010). Therefore, change in poverty over time inherits the decomposability of the growth incidence curve.

The term $\psi(y|z)$ in (2.29) is a convex and decreasing function measuring deprivation for an individual with a level of economic welfare equal to y[27]. A change in poverty associated with the growth pattern depicted by the incidence curve g(y) is given by:

$$\Delta P = \int_0^{m_y} yI(y \leq z)\psi'(y|z)g(y)dF(y) \quad (2.30)$$

Which using (2.19), can be equivalently expressed as:

$$\Delta P = \int_0^{m_y} I(y \leq z)y\psi'(y|z)g_X(y)dF(y) + \int_0^{m_y} I(y \leq z)y\psi'(y|z)g_S(y)dF(y) \quad (2.31)$$

The first term on the right represents the composition effect and the second term the structural effect. The detailed decomposition of the GIC carries over to variations in poverty outcomes that are based on additively separable poverty measures and RIF regression analysis can similarly be used. The relevant influence function is (Cowell and Victoria-Feser 1996, Essama-Nssah and Lambert 2012):

$$IF(y; P, F) = I(y \leq z)\psi(y|z) - P(F; z) \quad (2.32)$$

and the corresponding recentered influence function is:

$$RIF(y; P, F) = I(y \leq z)\psi(y|z) \quad (2.33)$$

For these poverty measures, one can also drop the linearity assumption and use nonlinear specifications of the RIF regression Fairlie (2005) proposes an extension of the Oaxaca-Blinder decomposition to logit and probit models while Bauer and Sinning (2008) explain how to extend the method to nonlinear models in general. Following these authors, we write the conditional mean outcome as $E(y_t|x_t; \beta_t), t = 0, 1$. The counterfactual mean outcome when endowments in period 1 are valued under the (reward) regime of period 0 is equal to the following: $E(y_1^c|x_1; \beta_0)$. The observed difference in mean outcomes can therefore be decomposed as follows.

$$\Delta_O^\mu = [E(y_1^c|x_1; \beta_0) - E(y_0|x_0; \beta_0)] + [E(y_1|x_1; \beta_1) - E(y_1^c|x_1; \beta_0)] \quad (2.34)$$

This expression is analogous to (2.8). The first term on the right is the composition effect and the second term is the structural effect. This is an aggregate

[27] For the Foster-Greer-Thorbecke (FGT) family of poverty measures, the deprivation function is $\psi(y|z) = \left(1 - \frac{y}{z}\right)^\alpha, \alpha \geq 0$. For the Watts measure, we have $\psi(y|z) = ln\left(\frac{z}{y}\right)$.

decomposition. The detailed decomposition is not as easy to obtain for nonlinear models as it is for linear models.

3. Empirical Results

To illustrate the power and usefulness of the methodology outlined above, we apply it to data for Cameroon for 2001-2007. Both the 2001 and 2007 household surveys are nationally representative with sample size 11,000 and 11,388 observations respectively. We use both OLS and RIF regression to link expenditure (in logs) to household characteristics, and we perform the relevant counterfactual decomposition to separate the endowment effect from the price effect for the growth incidence curve and also for variations in some social evaluation functions (for welfare and poverty). The discussion of the results focuses on the pattern of growth and its social impact.

3.1. The Pattern of Growth

We characterize the observed pattern of growth between 2001 and 2007 along three dimensions: (1) returns to selected household characteristics; (2) the relative importance of the composition and structural effects; and (3) the urban bias.

Returns to Selected Household Characteristics

We consider four broad categories of characteristics: (1) Demographics (gender of household head, age of household head, and household composition in terms of various age groups up to age 25); (2) Household and community assets (years of schooling of head of household, land ownership, access to credit, at least one migrant in household, distance to nearest hospital, distance to nearest tarred road); (3) Sector of employment (public sector, formal private sector, smallholder agriculture, informal non-agriculture, unemployed; and (4) Area/province of residence[28].

[28] Our choice of dummy variables implies that the reference household (conditional on characteristics represented by continuous variables) lives in the rural area of the central province and has a head who is female and out of the labor force, with no access to credit and no migrant in the household.

Our estimates of the marginal impact of each characteristic on household welfare in 2001 and 2007 are reported in tables A.1 and A.2 in the appendix. These tables show the coefficients and the associated standard errors for OLS and selected RIF regressions. We focus first on the OLS results. All demographic variables are statistically significant. As expected, an increase in any component of household membership reduces welfare. The male dummy variable has a negative sign in 2001 and a positive one in 2007. However, the 2007 coefficient is not significantly different from zero. Thus male-headed households do not necessary fare better than the reference female-headed households in either year, other things being equal.

Among the remaining non-geographical characteristics, the following have the highest positive and statistically significant impact on household welfare: (1) formal sector employment (public or private), (2) access to credit and (3) years of schooling of the head of household. Having at least one migrant in the household has no significant impact on welfare in either year. Similarly, land ownership does not make any difference, on average. The coefficient for agricultural employment is statistically significant in both years but has a negative sign. These regression results confirm that urban residence has a strong positive impact on welfare.

The OLS estimates give only average impacts for the characteristics under consideration. Results from RIF regressions will enable us in addition to appreciate the extent of heterogeneity in these impacts across quantiles which, as the reader will see, constitutes a considerable additional contribution to knowledge in respect of Cameroon. To keep our story manageable, we focus on three groups of covariates, namely household assets (education of head, access to credit, land ownership and having at least one migrant), sector of employment and area of residence (urban-rural). The effects of these characteristics are plotted against quantiles in figure A.1 in the appendix. Panel A shows plots of the unconditional quantile regression coefficients for education and access to credit. Returns to education (in terms of real per adult equivalent expenditure) are positive and statistically significant across all quantiles. Not surprisingly, the living standard increases with education over the whole distribution. With the exception of the lower end of the distribution and the segment between the 66th and 87th percentiles, the impact of education is significantly higher in 2001 than in 2007. This could be a manifestation of the lack of economic growth experienced by the country over that period. Indeed the lack of employment opportunities for the educated is a latent source of social tension in Cameroon.

The quantile plot for the returns to access to credit has an inverted U-shape in the low-end of the distribution (up to the 56th percentile) in 2001. Thus access to credit in 2001 increases inequality in the lower end of the distribution (up to the 25th percentile) and dampens inequality between the 25th and the 56th percentiles. In 2007, the effect has a U-shape in the same range of the distribution. The 2001 curve dominates the 2007 one. The effect of having access to credit is flat in the upper half of the distribution and there is no significant difference between the two years. Panel B shows the marginal impacts of land ownership and migration. Overall, land ownership has a very small positive impact on households in the low end of the distribution, particularly in 2007 and at the upper end of the distribution. Having at least one migrant in the household in 2001 made no significant difference for most households over the entire distribution except in the neighborhood of the 10th percentile where the impact is statistically significant and negative. In 2007, this factor has a positive impact very low in the distribution (up to the 10th percentile) and between the 60th and the 97th percentile. In addition, it contributes to decreasing inequality in the lower parts of those segments of the distribution in 2007.

The effects of formal sector employment are presented in panels C and D. Panel C reveals that returns to employment in the private sector are positive in 2001 for households located beyond the 12th percentile. In 2007, these returns are negative for most of that range up to the 75th percentile. For the lowest in the distribution, private sector employment brings positive returns only in 2007. The pattern of returns is similar for the public sector except that returns for 2007 are negative only over a very short range (from the 13th to the 24th percentile). A comparison of the two sectors in panel D shows that there is a reversal in the relative pattern of the returns to public and formal private sector employment between 2001 and 2007. The configuration of these quantile curves suggests two things. First, there is no advantage for the poor to be engaged in the formal sector. Second, the sluggish growth experienced between 2001 and 2007 may have hurt households engaged in the private sector more that those in the public sector.

We note from panel E that households engaged in agriculture are worse off across quantiles and years, than those employed in the other sectors of the economy[29]: the penalty associated with being engaged in agriculture hurts households at the lower end of the distribution more than those at the top. The marginal impact of urban residence in both years suggests that urban residence increases inequality in the low end of the distribution and decreases it in the top end. This pattern of unconditional quantile regression coefficients confirms that urban households are generally better off than their rural counterparts and that this urban bias has been increasing over time.

Composition versus Structure

Figure 3.1 shows a decomposition of the total variation in the distribution of log per capita expenditure (essentially, of the GIC) into two components. The first component is due to changes in the distribution of characteristics while the second represents the contribution of changes in the distribution of returns to those characteristics. Overall, the price (or structural) effect has a U-shape while the composition effect has an inverted U-shape. The structural effect dominates at the lowest end of the distribution while the composition effect dominates in the middle, from the 12th to the 76th percentile. The structural effect tends to decrease inequality at the lowest end of the distribution while the composition effect tends to increase it. The configuration of the three curves implies that the level of the GIC is determined mainly by the composition effect while the shape of the curve is explained mostly by the structural effect. In particular, the fact that people located at the bottom of the distribution up to the 10th percentile have experienced an income growth greater than average is due to the structural effect while the gains beyond that point are mainly due to the composition effect.

[29] The results for the informal sector, not shown, have the same pattern as those for smallholder agriculture.

Figure 3.1 A Decomposition of Growth Incidence in Cameroon (2001-2007)

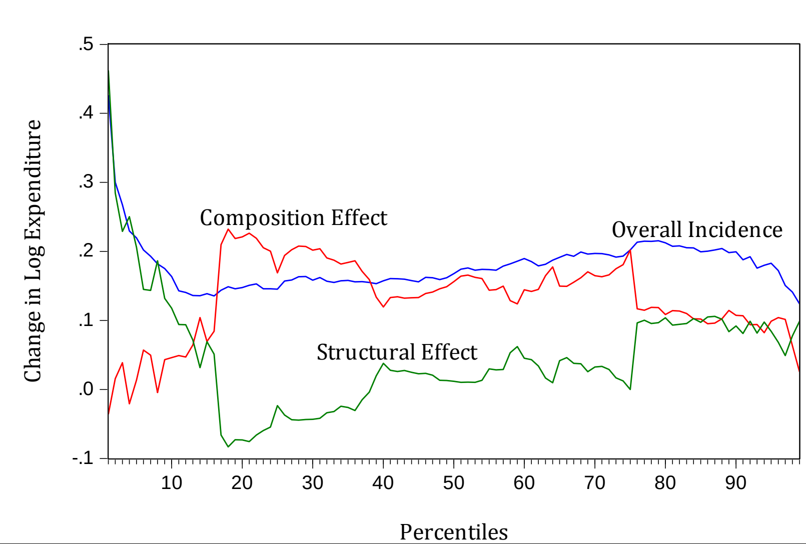

What are the factors driving both the composition and the structural effects? We disaggregate these two components on the basis of sets of covariates. The results are presented in figures 3.2 to 3.4. The left panel of figure 3.2 compares the full composition effect to the contribution of household demographics to this full effect. The right panel compares the same full effect and the contributions of household assets, sector of employment and geography. These results show that both the level and the dispersion of the full composition effect are mostly accounted for by household demographics.

Figure 3.2: Composition Effects

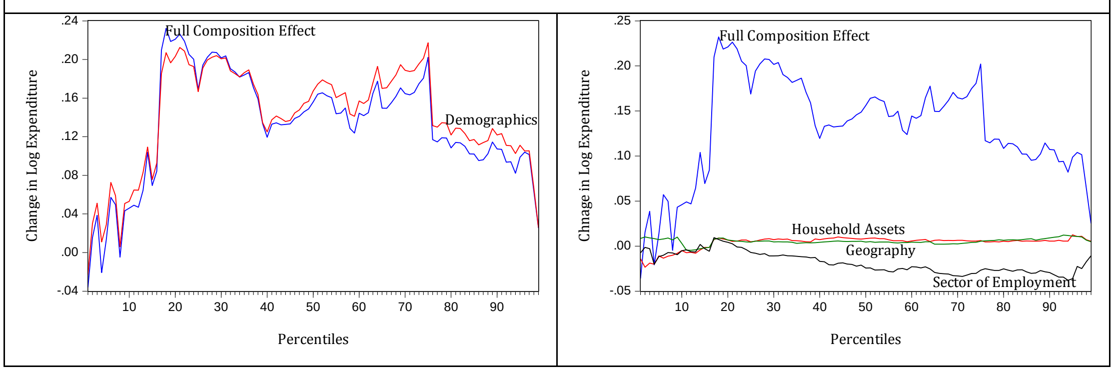

Figure 3.3 shows the contributions of various household assets to the composition effect. The figure reveals that the contribution of assets to the composition effect is mostly accounted for by changes in the distribution of years of schooling.

Figure 3.3: Composition Effects of Household Assets

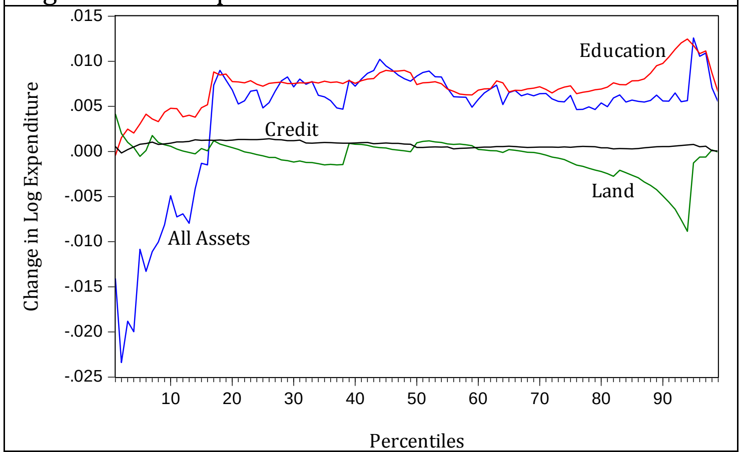

Finally, figure 3.4 presents results of a decomposition of the structural effect. These results suggest that the overall shape of the structural effect is determined by the sector of employment and, to a certain extent, geography. These are the two characteristics that might explain the negative values of the structural effect in some parts of the lower end of the distribution. This negative contribution is mitigated to some extent by household demographics.

Figure 3.4: Structural Effects

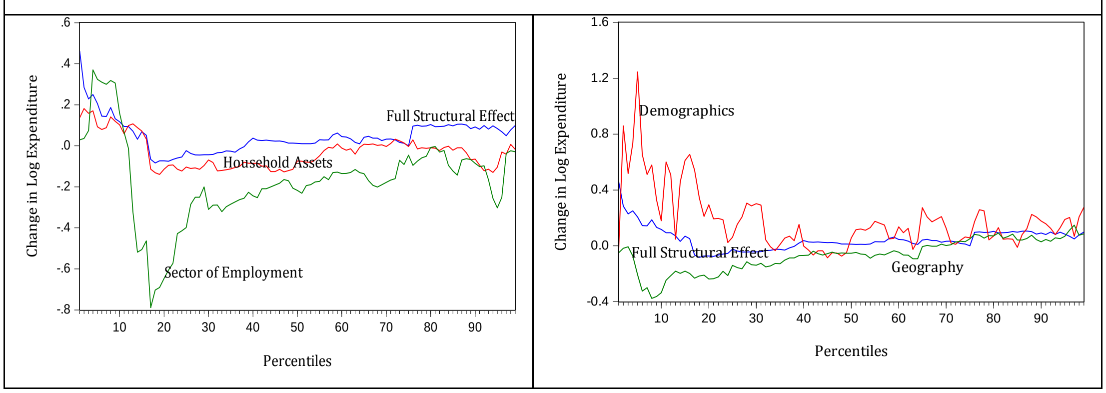

Urban Bias

The decomposition method discussed in this paper can be applied to the analysis of differences in distribution between groups such as rural versus urban. Figure 3.5 shows a comparison of the rural-urban differential in living standards in Cameroon for 2001 and 2007. The fact that the curve for 2007 dominates that for 2001 reveals that the gap between the rural and urban sector has been growing over time. In both years the total differential generally increases across quantiles, implying that dispersion increases at all points of the distribution. However, in 2001 the increase is steeper at the top end of the distribution, while in 2007 it is steeper at the low end. This observation implies that, in 2001 urban residence increased inequality more at the top of the distribution compared to the bottom. The opposite happened in 2007.

Figure 3.5 Change in Rural-Urban Differential

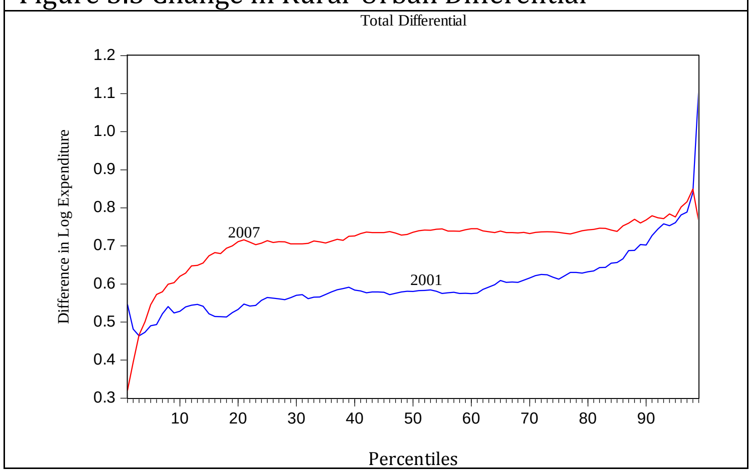

To further explore what may lie behind this configuration of *urban bias* in Cameroon, we use the same decomposition technique that we applied to the growth incidence curve above. Figure 3.6 shows the results. Focusing first on 2001, we notice that overall, the curve depicting the structural effect tends to follow a U-pattern while that representing the composition effect has, more or less, an inverted U-shape. Furthermore the structural effect dominates the composition effect over the whole range of the distribution, except between the 76th and the 88th quantiles. This shows that the greater increase in inequality at the top of the distribution in

2001 is due to the structural effect. The composition effect is pulling in the opposite direction.

Figure 3.6 A Decomposition of the Rural-Urban Differential

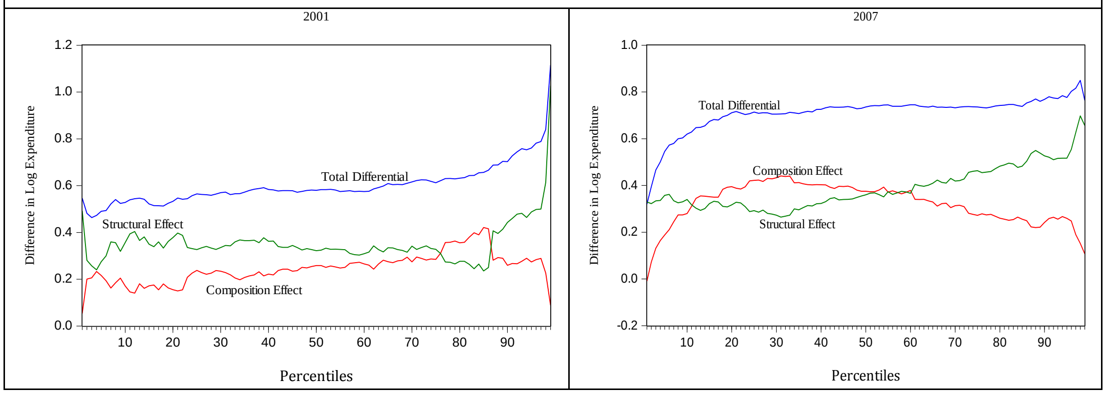

In 2007 the inverted U-shape of the composition effect is more pronounced than in 2001. The curve representing the structural effect has more or less the same shape as in 2001. It tends to fall until the 31st quantile then increases monotonically afterwards. The composition effect dominates the structural effect between the 11th quantile and the median. Both curves are very similar between the median and the 60th quantile. Past that point the structural effect clearly dominates the composition effect. Considering the overall profile of the total rural-urban differential, we note that the effect of urban residence on inequality is mostly driven by the composition effect at the low end of the distribution and by the structural effect (with some dampening by the composition effect) at the top end of the distribution.

Figure 3.7 shows how each of these two components evolved between 2001 and 2007. The key point that emerges from these observations is that composition contributes proportionately more to the increase in the rural-urban gap at lower quantiles while structure accounts for more of this increase at higher quantiles. In other words, differences in household characteristics matter more for the poorest households (particularly in 2007) than returns to those characteristics. The reverse is true for better-off households. This finding suggests that prevailing social arrangements treat the people at the bottom of the distribution alike whether they live in urban areas or not. At the top of the distribution social arrangements in urban areas reward better the set of characteristics than arrangements in the rural areas [30].

Figure 3.7 Changes in Composition and Structural Effects over Time

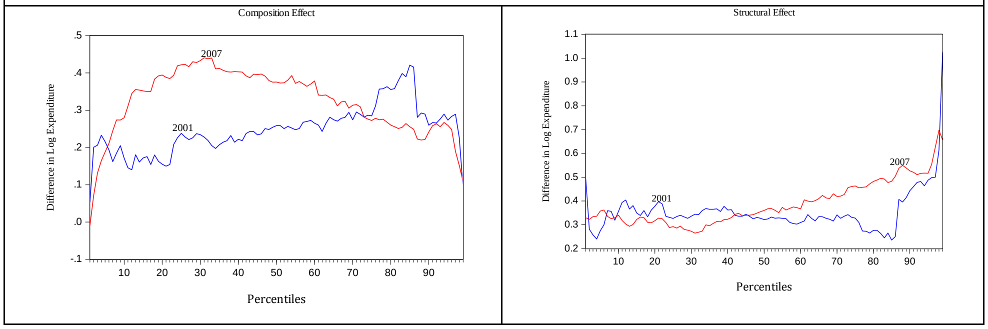

3.2. Social Impact

In figure 3.1, the curve representing the composition effect intersects the one representing the structural effect at several points: whether or not the composition effect dominates the structural effect depends on location on the GIC. A global assessment can however be made on the basis of one of the social evaluation functions discussed in section 2. First of all, we perform a standard Oaxaca-Blinder decomposition for the mean of the distribution of household expenditure, and we find that composition accounts for 73 percent of the overall variation while structure accounts for the rest (i.e. 27 percent).

Table 3.1 Accounting for Changes in the Atkinson Social Welfare Function

Aversion Level	Overall	Composition	Structure
Zero	100.00	66.16	33.84
One	100.00	72.86	27.14
Two	100.00	61.55	38.45

Source: Authors' calculation

Next, we decompose changes in the Atkinson social welfare function for three levels of inequality aversion, namely for $\varepsilon=0$, 1, and 2. (When epsilon=1 the decomposition agrees with the standard Oaxaca-Blinder one). The results are

[30] Nguyen et al. (2007) find a similar pattern for the case of Vietnam using a conditional quantile regression decomposition method proposed by Machado and Mata (2005). They explain their results by noting that, in accounting for rural-urban gap in well-being, one should not expect the structural effect to be important at the bottom of the distribution because poor people tend to work in jobs that pay little above the subsistence level. However, at the very top of the distribution, urban markets pay more for the same bundle of attributes than rural markets. One can therefore expect the structural effect to be more important than the composition effect in the upper end of the distribution. This explanation is also relevant to our case.

presented in table 3.1. When ε rises to two, the relative importance of the composition effect begins to decrease and that of the structural effect begins to increase.

Table 3.2 Accounting for Changes in the S-Gini Social Welfare Function

Aversion Level	Overall	Composition	Structure
One	100.00	72.86	27.14
Two	100.00	72.98	27.02
Three	100.00	66.92	33.08
Four	100.00	59.49	40.51
Five	100.00	52.25	47.75
Six	100.00	45.65	54.35

Source: Authors' calculation

Results for the decomposition of variations in the S-Gini social welfare function are presented in table 3.2. We let the aversion parameter, ν, vary from 1 to 6. This social welfare function agrees with the standard Oxaca-Blinder decomposition when ν=1. As ν increases beyond two, the importance of composition decreases relative to that of structure. When ν=6 the structure becomes more important than composition. This is consistent with our findings based on the decomposition of the GIC presented in figure 3.1.: structure dominates composition at the lower end of the distribution which the S-Gini evaluation function weighs more than the upper end in this case.

Table 3.3 Accounting for Changes in Poverty Outcomes

	Overall	Compo. 1	Struct. 1	Compo. 2	Struct. 2	Compo. 3	Struct. 3
FGT0	-0.0052	NA	NA	-0.0462	0.0410	-0.1244	0.1192
FGT1	-0.0055	-0.0407	0.0352	-0.0259	0.0204	-0.0537	0.0481
FGT2	-0.0056	-0.0171	0.0114	-0.0089	0.0033	-0.0218	0.0162
Watts	-0.0137	-0.0527	0.0390	-0.0325	0.0188	-0.0720	0.0583

Source: Authors' Calculations

On the basis of the available household data, we estimate the incidence of poverty to be 40 percent in 2001. Focusing on the segment of the distribution representing the poor (i.e. up to the 40^{th} percentile), we decompose the observed changes in poverty into composition and structural effects. Results are presented in table 3.3 for the headcount ratio (FGT0), the poverty gap (FGT1), the squared poverty gap (FGT2) and the Watts measure. The columns "Compo. 1" and "Struct. 1" show the composition and structural effects derived from the decomposition of the growth incidence curve. The columns "Compo. 2"and "Struct. 2" show results based on a linear RIF regression while those in the last two columns are based on nonlinear RIF regression using the logit model for the headcount ratio and the tobit model for the other poverty measures.

These results clearly show that, for the poor, the composition and the structural effects pull in opposite directions regardless of the poverty measure or the RIF regression model used. In all cases, the composition effect tends to reduce poverty while the structural effect tends to increase it. Overall, the composition effect dominates the structural effect for this segment of the distribution. This is consistent with the configuration of the growth incidence curve and the verdict rendered by social welfare functions of the second-order class for some choice of the aversion parameter. Thus the meager reduction in poverty observed in Cameroon between 2001 and 2007, is mainly due to the composition effect. This conclusion is robust with respect to the choice of poverty measures and RIF regression models. This robustness shows that the simple RIF regression approach, with its linear approximation, works pretty well. It produces results that are qualitatively similar to those based on more demanding non-linear specifications.

4. Concluding Remarks

For policymaking purposes, it is important to be able to identify the forces that shape a pattern of growth as revealed by the growth incidence curve. This paper demonstrates the use of RIF regression proposed by Firpo, Fortin and Lemieux (2009) to link the GIC to individual (or household) characteristics and to perform counterfactual decomposition *à la* Oaxaca-Blinder for the entire growth incidence curve following Fortin, Lemieux and Firpo (2011).

For generalized utilitarian social welfare functions and additively separable poverty measures, we demonstrate that variations in these functions are different aggregations of the informational content of the growth incidence curve and therefore inherit the decomposability of the GIC. A key implication is that one does not need to compute separate influence functions for the social evaluation functions considered here in order to decompose their variations *à la* Oaxaca-Blinder. The influence function for a quantile is enough for the job.

Furthermore, assuming the RIF regression to be linear leads to results that are qualitatively similar to those produced by non-linear specifications, and has the additional advantage of facilitating detailed decomposition that identifies the contribution of each covariate or set of covariates to the composition and structural effects not only for the GIC but for variations in the chosen social evaluation functions as well. This is a significant methodological finding that should give comfort to analysts who might be worried about the quality of the linear approximation underlying the simple RIF regression approach.

Our application of this methodology to data for Cameroon for 2001 and 2007 reveals a number of facts about the pattern of growth in that country and its social implications. The level of the growth incidence curve is driven by the composition (or endowment) effect while its shape is accounted for by the structural (or price) effect. The relationship between the two over the whole distribution indicates that, up to the 10^{th} quantile the observed gains are due to the structural effect. Beyond the 10^{th} quantile, gains are mainly due to the composition effect. The structural effect tends to reduce inequality at the lowest end of the distribution while the composition effect tends to increase it. Further disaggregation of these effects reveals that the composition effect is determined mainly by household demographics while the structural effect is accounted for mostly by the sector of employment and geography.

There is significant urban bias in the pattern of growth in Cameroon. Our results show that the gap between the rural and urban sectors has been growing over time. Overall, the effect of urban residence on inequality is mostly driven by the composition effect at the lower end of the distribution and by the structural effect at the top end of the distribution.

When the relative importance of the composition and structural effects is assessed on the basis of their social implications, we find that the composition effect accounts for the lion's share of the observed variation in social outcomes. In particular, this effect tends to reduce poverty while the structural effect tends to increase it. This conclusion is robust with respect to the choice of poverty measures and RIF regression models. Thus, for the poor, differences in household characteristics matter more than returns to those characteristics. This suggests that prevailing social arrangements do not afford the same opportunities for the poor and the non-poor. Our empirical findings also suggest that, when dominance fails, the S-Gini social welfare function might be a more distribution-sensitive criterion than the poverty measures considered here. This social evaluation function might therefore be more suitable for assessing the extent to which a pattern of growth is inclusive.

What is the policymaker to make of these findings? The goal of the Poverty Reduction Strategy adopted in 2003 was to reduce poverty by promoting inclusive growth through improved governance and service delivery, and by letting the private sector drive the growth process. Subject to data reliability, our findings suggest that four years into the implementation of the strategy there has been no progress towards achieving that goal. What explains this outcome and what can be done about it?

The potential impact of government intervention on the level and pattern of economic growth depends on how such intervention affects factor *accumulation* and *productivity*. The composition effect is tied to accumulation while the structural effect is an indicator of productivity. The former turned out to be more important for the poor, inequality enhancing and mostly accounted for by the distribution of years of schooling. Thus, there is a need to further study why the delivery of educational services is not working for the poor. The pattern of the structural effect suggests that the weak growth performance observed over the period under study maybe due to low factor productivity. This in turn may be due to poor quality of institutions. Finally, the increased urban bias and the fact that households engaged in agriculture are worse off across quantiles and over time suggest either neglect of agriculture, (a vital sector of the economy) or that investments in that sector, if any, may not be productive. Ultimately, the policymakers in Cameroon and their partners need to stop and ponder whether, all along, they have been doing the right things the right way.

Appendix: Returns to Household Characteristics

Table A.1:
OLS and RIF Regression Coefficients on Log Expenditure, 2001

Eq Name:	OLS	Quantile 10	Quantile 25	Quantile 50	Quantile 75	Quantile 90
Dep. Var:	LPCEXP	RIFQT_10	RIFQT_25	RIFQT_50	RIFQT_75	RIFQT_90
Constant	13.159	12.210	12.572	12.924	13.729	13.828
	(0.064)**	(0.161)**	(0.111)**	(0.079)**	(0.086)**	(0.128)**
Male	-0.174	-0.059	-0.192	-0.206	-0.269	-0.171
	(0.015)**	(0.038)	(0.026)**	(0.018)**	(0.020)**	(0.030)**
Age of Head	-0.011	0.007	-0.005	-0.008	-0.013	-0.009
	(0.002)**	(0.005)	(0.004)	(0.003)**	(0.003)**	(0.004)*
Age Head Squared	0.0001	-0.0001	0.00001	0.0001	0.0001	0.0001
	(0.000)**	(0.0001)*	(0.000)	(0.000)*	(0.000)**	(0.000)*
Age<5 (% of Household)	-0.008	-0.008	-0.009	-0.006	-0.008	-0.010
	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 5 to <10 (%HH)	-0.012	-0.015	-0.015	-0.011	-0.013	-0.014
	(0.0004)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 10 to < 15 (%HH)	-0.009	-0.008	-0.012	-0.010	-0.011	-0.012
	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 15 to <20 (%HH)	-0.008	-0.006	-0.010	-0.007	-0.008	-0.009
	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 20 to <25 (%HH)	-0.006	-0.008	-0.008	-0.005	-0.006	-0.007
	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Schooling (Years)	0.035	0.024	0.036	0.037	0.036	0.049
	(0.002)**	(0.004)**	(0.003)**	(0.002)**	(0.002)**	(0.003)**
Land	0.00004	-0.0003	0.0002	-0.001	0.001	0.003
	(0.0002)	(0.0005)	(0.0003)	(0.0002)*	(0.0002)**	(0.0004)**
Access to Credit	0.273	0.345	0.502	0.168	0.174	0.203
	(0.018)**	(0.047)**	(0.032)**	(0.023)**	(0.025)**	(0.037)**
Has Migrant (s)	-0.003	-0.093	0.036	-0.009	-0.008	-0.014
	(0.011)	(0.028)**	(0.019)	(0.014)	(0.015)	(0.022)
Distance to Nearest Hospital	-0.008	-0.019	-0.005	-0.003	-0.001	-0.0003
	(0.001)**	(0.002)**	(0.002)**	(0.001)*	(0.001)	(0.002)
Distance to Nearest Tarred Road	0.001	0.003	-0.0002	0.002	0.0005	0.001
	(0.0002)**	(0.0004)**	(0.0003)	(0.0002)**	(0.0002)*	(0.0003)
Public Sector	0.137	-0.262	0.281	0.159	0.165	0.327
	(0.024)**	(0.062)**	(0.042)**	(0.031)**	(0.033)**	(0.043)**
Private Sector (Formal)	0.268	-0.061	0.393	0.267	0.271	0.346
	(0.025)**	(0.063)	(0.043)**	(0.031)**	(0.036)**	(0.049)**
Agriculture	-0.047	-0.241	0.205	-0.079	-0.156	-0.047
	(0.018)*	(0.046)**	(0.031)**	(0.023)**	(0.025)**	(0.036)
Non-Agriculture Informal	-0.003	-0.098	0.169	-0.053	-0.063	-0.026
	(0.020)	(0.050)	(0.035)**	(0.025)*	(0.027)*	(0.040)
Unemployed	0.075	-0.096	0.369	0.163	0.251	-0.256
	(0.037)*	(0.092)	(0.063)**	(0.045)**	(0.049)**	(0.073)**
Urban	0.358	0.327	0.411	0.425	0.324	0.367
	(0.017)**	(0.043)**	(0.030)**	(0.022)**	(0.023)**	(0.034)**

Adamaoua	0.0703	0.056	-0.012	0.119	0.036	0.012
	(0.032)*	(0.081)	(0.056)	(0.040)**	(0.043)	(0.064)
East	-0.021	-0.209	-0.159	-0.025	-0.011	-0.038
	(0.026)	(0.066)**	(0.046)**	(0.033)	(0.036)	(0.053)
Far-North	0.123	0.072	0.035	0.153	0.141	0.151
	(0.019)**	(0.049)	(0.034)	(0.024)**	(0.027)**	(0.039)**
Coast	0.110	0.052	0.033	0.105	0.143	0.103
	(0.023)**	(0.057)	(0.039)	(0.028)**	(0.031)**	(0.045)*
North	0.171	0.148	0.013	0.116	0.156	0.255
	(0.024)**	(0.061)*	(0.042)	(0.030)**	(0.033)**	(0.048)**
North-West	-0.097	-0.273	-0.206	-0.041	-0.032	-0.018
	(0.020)**	(0.051)**	(0.035)**	(0.025)	(0.027)	(0.040)
West	0.095	0.422	0.038	0.089	0.011	-0.048
	(0.019)**	(0.049)**	(0.034)	(0.024)**	(0.026)	(0.039)
South	0.118	0.391	0.348	0.291	-0.166	-0.140
	(0.036)**	(0.090)**	(0.062)**	(0.045)**	(0.048)**	(0.071)
South-West	0.040	0.027	0.043	0.149	-0.124	-0.125
	(0.022)	(0.056)	(0.039)	(0.027)**	(0.029)**	(0.044)**
Observations:	11000	11000	11000	11000	11000	11000
R-squared/Pseudo R2 :	0.365	0.097	0.199	0.291	0.256	0.156
F-statistic:	218.321	40.628	94.074	155.160	129.847	69.799

Source: Authors' Calculations (Standard Errors in Parentheses)

Table A.2: OLS and RIF Regression Coefficients on Log Expenditure, 2007

Eq Name:	OLS	Quantile 10	Quantile 25	Quantile 50	Quantile 75	Quantile 90
Dep. Var:	LPCEXP	RIFQT_10	RIFQT_25	RIFQT_50	RIFQT_75	RIFQT_90
Constant	13.304	12.197	13.121	13.238	13.719	13.899
	(0.063)**	(0.141)**	(0.111)**	(0.088)**	(0.102)**	(0.142)**
Male	0.070	-0.155	-0.225	0.0714	0.402	0.244
	(0.037)	(0.082)	(0.065)**	(0.052)	(0.059)**	(0.083)**
Age of Head	-0.013	-0.008	-0.017	-0.011	-0.014	-0.005
	(0.002)**	(0.004)*	(0.003)**	(0.003)**	(0.003)**	(0.004)
Age Head Squared	0.0001	0.0001	0.0001	0.0001	0.0001	0.00004
	(0.000)**	(0.000)**	(0.000)**	(0.000)**	(0.000)**	(0.000)
Age<5 (% of Household)	-0.005	0.0003	-0.006	-0.006	-0.008	-0.010
	(0.0004)**	(0.001)	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 5 to <10 (%HH)	-0.008	-0.004	-0.008	-0.008	-0.012	-0.012
	(0.0004)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 10 to < 15 (%HH)	-0.008	-0.011	-0.009	-0.007	-0.007	-0.008
	(0.0004)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 15 to <20 (%HH)	-0.007	-0.005	-0.006	-0.006	-0.009	-0.010
	(0.0004)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)**
Age 20 to <25 (%HH)	-0.001	0.0004	-0.0003	-0.001	-0.002	-0.002
	(0.0004)**	(0.001)	(0.001)	(0.001)	(0.001)**	(0.001)**
Schooling (Years)	0.028	0.022	0.021	0.028	0.0311	0.038
	(0.001)**	(0.003)**	(0.003)**	(0.002)**	(0.002)**	(0.003)**
Land	0.001	0.001	0.001	0.0001	0.0003	0.003
	(0.0003)**	(0.001)	(0.001)**	(0.0004)	(0.001)	(0.001)**
Access to Credit	0.122	0.326	-0.003	0.153	0.183	0.188
	(0.016)**	(0.037)**	(0.029)	(0.023)**	(0.027)**	(0.037)**
Has Migrant (s)	0.007	0.059	-0.047	-0.027	0.053	0.025
	(0.009)	(0.021)**	(0.016)**	(0.013)*	(0.015)**	(0.021)

Distance to Nearest Hospital	-0.002	0.003	-0.003	-0.003	0.0004	-0.002
	(0.001)**	(0.001)**	(0.001)**	(0.001)**	(0.001)	(0.001)
Distance to Nearest Tarred Road	0.001	0.001	0.001	0.0001	0.0005	-0.0003
	(0.0001)**	(0.0003)**	(0.0002)**	(0.0002)	(0.0002)*	(0.0003)
Public Sector	0.101	0.146	0.018	-0.034	0.294	0.314
	(0.039)*	(0.088)	(0.070)	(0.055)	(0.064)**	(0.089)**
Private Sector Formal	0.011	0.121	0.018	-0.157	0.014	0.244
	(0.040)	(0.091)	(0.072)	(0.057)**	(0.066)	(0.091)**
Agriculture	-0.253	-0.143	-0.363	-0.348	-0.149	-0.091
	(0.036)**	(0.080)	(0.063)**	(0.050)**	(0.058)**	(0.081)
Non-Agriculture Informal	-0.145	0.127	-0.051	-0.182	-0.160	-0.221
	(0.035)**	(0.078)	(0.062)	(0.049)**	(0.057)**	(0.079)**
Unemployed	-0.191	-0.027	-0.142	-0.232	-0.135	-0.072
	(0.040)**	(0.090)	(0.071)*	(0.057)**	(0.065)*	(0.091)
Urban	0.429	0.076	0.305	0.549	0.683	0.564
	(0.016)**	(0.036)*	(0.029)**	(0.023)**	(0.026)**	(0.037)**
Adamaoua	0.012	-0.054	0.002	-0.067	0.083	0.124
	(0.023)	(0.051)	(0.041)	(0.033)*	(0.038)*	(0.053)*
East	-0.159	-0.413	-0.303	-0.149	-0.106	-0.029
	(0.024)**	(0.054)**	(0.042)**	(0.034)**	(0.039)**	(0.054)
Far-North	-0.177	-0.744	-0.449	-0.088	0.002	0.116
	(0.017)**	(0.038)**	(0.030)**	(0.024)**	(0.028)	(0.038)**
Coast	-0.014	0.085	0.149	-0.025	-0.251	-0.208
	(0.034)	(0.076)	(0.059)*	(0.047)	(0.055)**	(0.076)**
North	-0.138	-0.233	-0.398	-0.158	-0.001	0.079
	(0.018)**	(0.041)**	(0.032)**	(0.026)**	(0.029)	(0.041)
North-West	-0.112	-0.403	-0.311	-0.096	-0.022	0.059
	(0.021)**	(0.046)**	(0.037)**	(0.029)**	(0.034)	(0.047)
West	0.104	0.079	0.223	0.154	-0.001	-0.078
	(0.020)**	(0.045)	(0.035)**	(0.028)**	(0.032)	(0.045)
South	0.096	0.032	0.203	0.175	0.096	-0.113
	(0.029)**	(0.066)	(0.052)**	(0.041)**	(0.048)*	(0.066)
South-West	0.149	0.147	0.192	0.264	0.025	0.043
	(0.020)**	(0.045)**	(0.036)**	(0.028)**	(0.033)	(0.046)
Observations:	11388	11388	11388	11388	11388	11388
R-squared/Pseudo R2:	0.499	0.199	0.314	0.380	0.332	0.174
F-statistic:	390.739	97.382	179.013	240.464	194.267	82.759

Source: Authors' Calculations (Standard Errors in Parentheses)

Figure A.1: Returns to Selected household characteristics

Panel A: Returns to Education and Access to Credit

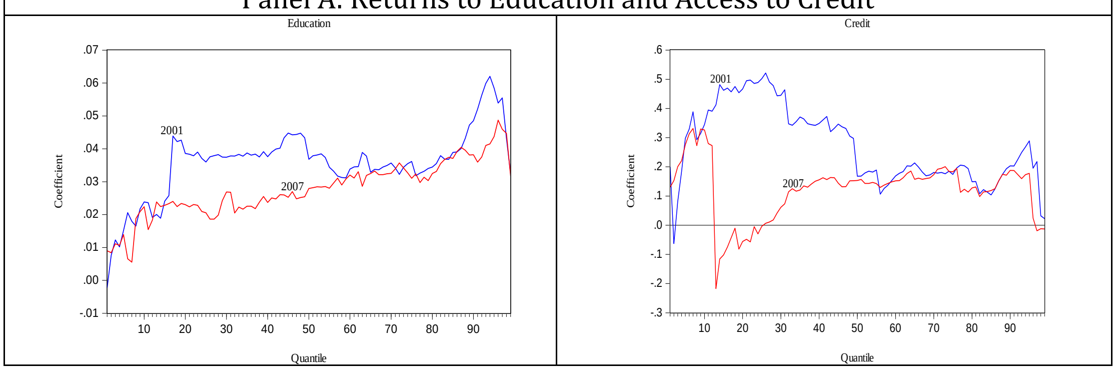

Panel B: Returns to Land and Migration

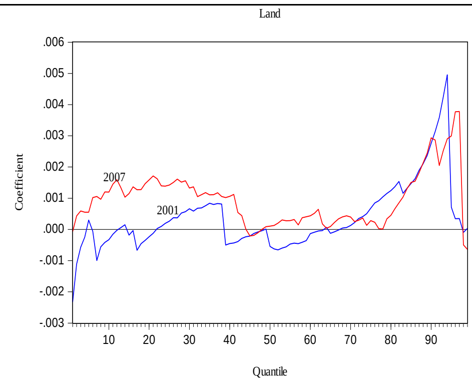

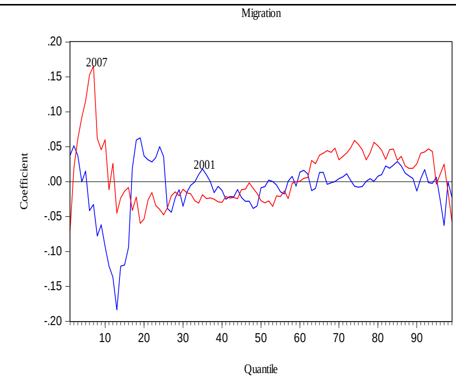

Panel C: Returns to Formal Sector Employment

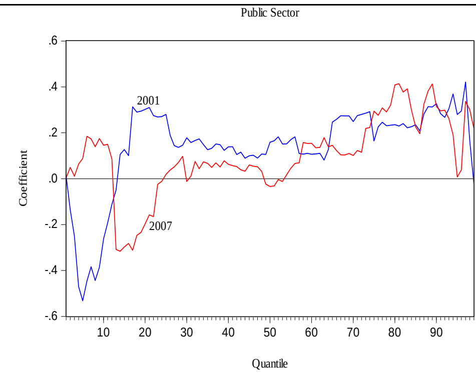

Panel D: Private and Public Sectors Compared

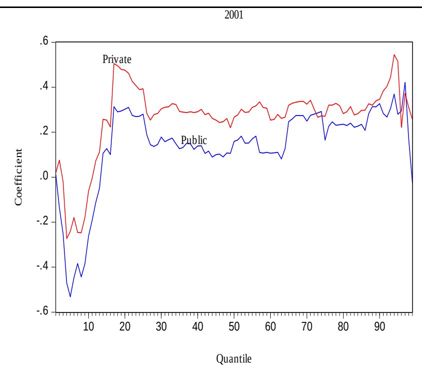

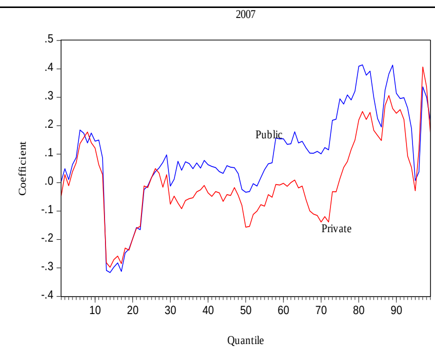

Panel E: Returns to Smallholder Agriculture and Urban Residence

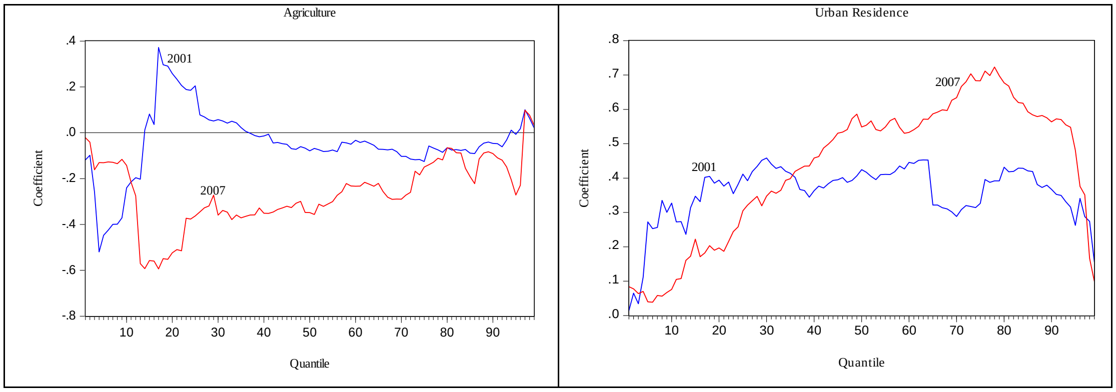

References

Atkinson Anthony B. 1970. On the Measurement of Inequality. Journal of Economic Theory, **2,** 244-263.

Bauer Thomas K. and Sinning Mattias. 2008. An Extension of the Blinder –Oaxaca Decomposition to Nonlinear Models. Advances in Statistical Analysis, Vol. 92, No. 2:197-2006.

Bitler Marianne P., Gelbach Jonah B., and Hoynes Hilary W. 2006. What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments. American Economic Review, Volume 96, No. 4, 988-1012.

Blandford David, Friedman Deborah, Lynch Sarah, Mukherjee Natasha and Sahn David E. 1994. Oil Boom and Bust: The Harsh Realities of Adjustment in Camerroon. In David E. Sahn (ed), *Adjusting to Policy Failure in African Economies*. Ithaca: Cornel University Press.

Bourguignon François. 2010. Non-anonymous Growth Incidence Curves, Income Mobility and Social Welfare Dominance. Journal of Economic Inequality DOI 10.1007/s10888-010-9159-7.

Bourguignon François, Ferreira Francisco H. G. and Leite Phillippe G. 2008. Beyond Oaxaca-Blinder : Accounting for Differences in Households Income Distributions. Journal of Economic Inequality, Vol. 6, No.2, 117-148.

Bourguignon François, Ferreira Francisco H. G. 2005. Decomposing Changes in the Distribution of Household Incomes: Methodological Aspects. In: Bourguignon, François, Ferreira Francisco H.G. and Lustig Nora (eds), The Microeconomics of Income Distribution Dynamics in East Asia and Latin America, pp. 17-46 Washington, D.C.: The World Bank.

Carneiro Pedro, Hansen Karsten T. and Heckman James J. 2002. Removing the Veil of Ignorance in Assessing the Distributional Impacts of Social Policies. Institute for the Study of Labor (IZA) Discussion Paper No. 453.

Commission on Growth and Development. 2008. The Growth Report: Strategies for Sustained Growth and Inclusive Development. Washington, D.C.: The World Bank.

Cowell Frank A. and Victoria-Feser Maria-Pia. 1996. Poverty Measurement with Contaminated Data: A robust Approach. European Economic Review, 40: 1761-1771.

Dadson Nick. 2012. Recentered Influence Functions and the Shapley Detailed Decomposition: An Application to Poverty in the United States. Mimeo, Simon Fraser University.

Drazen Allan. 2000. Political Economy in Macroeconomics. Princeton: Princeton University Press.

Duclos Jean-Yves and Araar Abdelkrim. 2006. Poverty and Equity: Measurement, Policy and Estimation with DAD. Springer, New York.

Essama-Nssah B. 2012. Identification of Sources of Variation in Poverty Outcomes. World Bank Policy Research Working Paper No. 5954 Wasington, DC.

Essama-Nssah B. 2005. A Unified Framework for Pro-Poor Growth Analysis. Economics Letters, **89**:216-221.

Essama-Nssah B. 2002. Assessing the Distributional Impact of Public Policy. World Bank Policy Research Working Paper No. 2883. The World Bank, Washington, D.C.

Essama-Nssah and Lambert Peter J. 2012. Influence Functions for Policy Impact Analysis, Chapter 6 in John A. Bishop, Rafael Salas (eds.) Inequality, Mobility and Segregation: Essays in Honor of Jacques Silber (Research on Economic Inequality, Volume 20), Emerald Group Publishing Limited, pp. 135 - 159

Essama-Nssah B., Bassole Leandre and Saumik Paul. 2010. Accounting for Heterogenity in Growth Incidence in Cameroon. World Bank Policy Research Working Paper No. 5464. Wasington, DC.

Essama-Nssah, B. and Lambert Peter J. 2009. Measuring Pro-Poorness: a Unifying Approach with New Results. Review of Income and Wealth, Series 55, No.3: 752-778.

Fairlie Robert W. 2005. An Extension of the Blinder-Oaxaca Decomposition Technique to Logit and Probit Models. Journal of Economic and Social Measurement, Vol. 30, No. 4: 305-316.

Ferreira Francisco H.G. 2010. Distributions in Motion: Economic Growth, Inequality and Poverty Dynamics. World Bank Policy Research Working Paper No. 5424. The World Bank, Washington, D.C.

Firpo Sergio, Fortin Nicole and Lemieux Thomas. 2009. Unconditional Quantile Regressions. Econometrica, Volume 77, No. 3: 953-973.

Fortin Nicole, Lemieux Thomas and Firpo Sergio. 2011. Decomposition Methods in Economics. In: Orley Ashenfelter and David Card (eds) Handbook of Labor Economics, Vol. 4A , pp. 1-102. Northolland, Amsterdam.

Foster James E. and Shorrocks Anthony F.: Subgroup Consistent Poverty Indices. Econometrica, Vol. 59, No.3, 687-709 (1991)

Foster James E., Greer Joel and Thorbecke Erik. 2010. The Foster-Greer-Thorbecke (FGT) Poverty Measures: 25 Years Later. Journal of Economic Inequality, 8: 491-524.

Foster James E., Greer Joel and Thorbecke Erik. 1984. A Class of Decomposable Poverty Measures. Econometrica, Vol. 52, No3, 761-766.

Government of Cameroon. 2003. Poverty Reduction Strategy Paper. Yaoundé: Republic of Cameroon.

Grimm, Michael. 2007. Removing the Anonymity Axiom in Assessing Pro-poor Growth. Journal of Economic Inequality, vol. 5, pp. 179–197.

Grosse Melanie, Harttgen Kenneth and Klasen Stephan. 2008. Measuring Pro-Poor Growth in Non-Income Dimensions. World Development, Vol. 36, No.6, 1021-1047.

Hampel F.R. 1974. The Influence Curve and its Role in Robust Estimation. Journal of the American Statistical Association, 60: 383-393.

Institut National de la Statistique (National Statistical Office). 2008. Troisième Enquête Camerounaise auprès des Ménages (ECAM3) : Tendances, Profil et Déterminants de la Pauvreté au Cameroun entre 2001-2007. Yaoundé : République du Cameroun.

Machado J.A.F. and Mata J. 2005. Counterfactual Decomposition of Changes in Wage Distributions Using Quantile Regression. Journal of Applied Econometrics, Vol. 20, No 4, 445-465.

Milgrom Paul. 2004. Putting Auction Theory to Work. Cambridge: Cambridge University Press.

Moulin, Hervé J.: Fair Division and Collective Welfare. The MIT Press, Cambridge (Massachusetts) (2003)

Nguyen Binh T., Albrecht James W., Vroman Susan B. and Westbrook M. Daniel. 2007. A Quantile Regression Decomposition of Urban-Rural Inequality in Vietnam. Journal of Development Economics **83**, 466-490.

Ravallion Martin: Growth, Inequality and Poverty. 2001. Looking Beyond Averages. World Development, Vol 29, No. 11, 1803-1815.

Ravallion Martin and Chen Shaohua. 2003. Measuring Pro-Poor Growth. Economics Letters **78**, 93-99.

Rothe Christoph. 2010. Decomposing Counterfactual Distributions. Toulouse School of Economics, mimeo.

Sen Amartya. 1995. Inequality Reexamined. Cambridge (Massachusetts): Harvard University Press.

Sen Amartya. 1989. The Concept of Development. In: Chenery Hollis and Srinivasan T.N. (eds) Handbook of Development Economics, Vol. 1, pp. 10-26.

Von Mises R. 1947. On the Asymptotic Distribution of Differentiable Statistical Functions. The Annals of Mathematical Statistics, **18**, 309-348.

Watts Harold. 1968. An Economic Definition of Poverty. In: Moynihan Daniel P. (ed.), On Understanding Poverty: Perspectives from the Social Sciences, pp. 316-329. Basic Books, New York.

Wilcox Rand R. 2005. Introduction to Robust Estimation and Hypothesis Testing (Second Edition). Elsevier, Amsterdam.

Yitzhaki Shlomo. 1983. On an Extension of the Gini Index of Inequality. International Economic Review, Vol. 24, No. 3, 617-628.

Zheng Buhong. 2000a. Poverty Orderings. Journal of Economic Surveys, Vol.14, No.4, 427-466.

Zheng Buhong. 2000b. Minimum-Distribution-Sensitivity, Poverty Aversion, and Poverty Orderings. Journal of Economic Theory, **95**, 116-137.

Zheng Buhong. 1993. An Axiomatic Characterization of the Watts Poverty Index. Economics Letters, **42**, 81-86.