

The Center Cannot Hold: Patterns of Polarization in Nigeria*

F. Clementi

University of Macerata, Macerata, Italy

A. L. Dabalén and V. Molini[†]

World Bank, Washington DC, USA

F. Schettino

Second University of Naples, Naples, Italy

June 19, 2014

*The paper enjoyed financial contribution from the World Bank. The authors would like to thank Federica Alfani (FAO - Food and Agriculture Organization of the United Nations), Rose Mungai (World Bank - Poverty Reduction and Economic Management Unit), and Ayago E. Wambile (World Bank - Poverty Reduction and Economic Management Unit) for excellent assistance with data preparation. Completion of this work would have not been possible without them. Of course, we are the sole responsible for all possible errors the paper may contain.

[†]Corresponding author: vmolini@worldbank.org.

Abstract

Recent analysis on consumption patterns in Nigeria seems to suggest an increase in inequality that could have offset the poverty-reducing benefits from sustained growth. Inequality increase is however just one aspect of the whole problem. Our hypothesis is that Nigeria is also going through a process of economic polarization. Broadly speaking, the notion of polarization is concerned with the disappearance or—as in the case of Nigeria—non-consolidation of the middle class, which occurs when there is a tendency to concentrate in the tails, rather than the middle, of the income/consumption distribution. The aim of this paper is to document the increasing polarization, paying a special attention to certain areas of the country, notably those where in the last decade economic growth had been stagnant or particularly non-inclusive with vast negative repercussions on social stability. An analysis of this type is rather new for Nigeria. The limited availability of comparable data has hindered an investigation that requires data series not too close in time: the process of polarization is generally slow and significant changes can be detected over long periods. The present paper tries to overcome this limitation by making use of recently developed survey-to-survey imputation techniques. To explore polarization, our study uses instead the relative distribution methodology ([Handcock and Morris, 1998, 1999](#)). This flexible and straightforward method provides a non-parametric framework for comparing the income (or other) distributions of two populations—either cross-sectional or over time—in a way that permits consideration of differences throughout the entire income range. Findings confirm the sharp increase of polarization. Compared to 2003, the 2013 consumption distribution is more concentrated in upper and lower deciles, while the middle deciles are progressively emptying out. A between-group analysis based on the six geo-political zones of the country also shows the emergence of a macro-regional gap: in fact, while the South South and South West regions contribute mainly to polarization in the upper tail of the national consumption distribution, households in the North East and North West zones—the conflict-stricken areas—are more likely to fall in the lower national deciles compared to the rest of the country.

JEL classification: C14; D31; D63; I32; Q34

Keywords: Nigeria; consumption expenditure; poverty and inequality; relative distribution; polarization

1 Introduction

Despite a stable and sustained growth, according to official numbers poverty reduction in Nigeria has not been up to general expectations. Poverty seems to have declined faster in the coastal South and around the Federal Capital, Abuja. Not every state improved. To the contrary, a large belt of north-eastern states have experienced a significant increase in poverty. The lack of a faster reduction in poverty despite a significant growth in GDP may be due to a fast increase in inequality (World Bank, 2013).

An increase in inequality is, however, just one aspect of the whole problem. Our hypothesis is that Nigeria is also undergoing through a process of increasing income polarization. Whereas inequality is the overall dispersion of the distribution, referring to the distance of every individual from the median or mean income, polarization is the combination of divergence from global and convergence on local mean incomes.

In income-polarized societies, people cluster around group means and tend to be far from the mean/median of the overall distribution. Within each group there is income homogeneity and often reducing income inequality: we can talk, thus, of “increasing identification”. Between the two groups, instead, we talk of “increasing alienation” (Duclos *et al.*, 2004). The combined effect of alienation and identification forces between two significantly sized groups leads to effective opposition, a situation that might give rise to social conflicts and tensions (Esteban and Ray, 1999, 2008, 2011). Also, the group at the top of the distribution has voice while the other group, made by those at the bottom, are voiceless in matters that affect their welfare and the society at large.

Another important aspect of the income polarization analysis is that it is concerned with the disappearance or—as in the case of Nigeria—non-consolidation of the middle class. This precisely occurs when in a society there is a tendency to concentrate in the tails, rather than the middle, of the income distribution. A well-off middle class is important to every society because it contributes significantly to economic growth, as well as to social and political stability (Easterly, 2001; Pressman, 2007). Also, the middle class constitutes the backbone of democracy (Birdsall, 2010) and is a key ingredient in guaranteeing sustainable economic growth and poverty reduction efforts in the long-term.

Nigeria represents an interesting case for undertaking a polarization analysis. As mentioned before, GDP and per capita income have steadily grown in the last decade, after GDP re-basing

Nigeria is likely to become the biggest African economy and yet clear signs of consolidation of a national middle class are limited. Moreover, the country is increasingly affected by sub-regional conflicts driven to a large extent by disaffected (alienated) groups.

Studies on polarization in Nigeria are few and have approached the relevant issues in a restrictive way. The limited attention paid to long-run patterns could have been due to data problems. [Aigbokhan \(2000\)](#) used the [Wolfson \(1994\)](#) polarization index and provided estimates for the country's urban and rural areas under national, male-headed, female-headed, and zones dimensions. They found that there was a higher degree of polarization in the rural areas in the 1990s, and while polarization increased in the country between 1985 and 1992, it declined in the rural areas, which is in contrast with the general belief of increased polarization.

[Araar \(2008\)](#) analyzed the 2003/2004 round of the National Living Standard Survey (NLSS), and tried to identify the main drivers behind polarization in the 1990s by comparing Nigeria to China. Using the same set of data, [Awoyemi and Araar \(2009\)](#) decomposed the Duclos-Esteban-Ray (DER) index of polarization ([Duclos et al., 2004](#)). Main results indicate a clear prevalence of the identification component vs. alienation, making authors to hypothesize that the existence of an ongoing polarization process. Also, they identify as main drivers of polarization the increasing divide between macro-zones, education and levels of occupation. Urban/rural divide is found to be insignificant. [Awoyemi et al. \(2010\)](#) extended the analysis looking at the polarization dynamics over the a longer time span, 1996–2004. Using data from two different household surveys, they find a reduction in polarization using both the [Foster-Wolfson \(1992\)](#) and DER index: from 0.30 to 0.25 and from 0.44 to 0.38, respectively. They also show that in the southern macro areas (South East and South West) indexes do not vary significantly. [Ogunyemi and Oni \(2011\)](#) and [Ogunyemi et al. \(2011\)](#) calculated the same indexes on households in rural areas only from 1980 to 2004, finding a similarly decreasing trend. More recent studies like [Ogunyemi \(2013\)](#) indicate a general invariance, with a limited tendency towards increase, by comparing the 2003/2004 round of the NLSS and the 2009/2010 round of the Harmonized National Living Standard Survey (HNLSS).

The present paper is innovative in several aspects. First, rather than just computing and comparing polarization indexes, we use a non-parametric framework (the “relative distribution” introduced by [Handcock and Morris, 1998, 1999](#)) and compare income throughout the entire income range. The relative distribution analysis requires at least two comparable survey rounds in order to investigate changes along the entire distribution. Since the lack of comparable

surveys has limited the scope of previous work, we use survey-to-survey techniques to produce two fully comparable distributions—and this can be regarded as the second aspect of novelty of the present study. Finally, the flexibility of the relative distribution tool allows an accurate analysis at macro-regional level too. Differently from previous contributions, another goal of this paper is to document also sub-national patterns of polarization. Nigeria is highly heterogeneous, so that drivers of polarization can indeed differ across macro regions. It is also worth mentioning that this focus on macro-regions is aimed at preparing the ground for future research on the link between polarization and regional conflicts.

Besides the introduction, the paper articulates in four additional sections. Section 2 presents the data and discusses the imputation strategy we use to obtain comparable data on household consumption. Section 3 reviews the approaches to measuring economic polarization and outlines the distinctive features of the relative distribution method. Section 4 details the main findings of the study. Section 5 concludes.

2 Data and Empirical Strategy

The comparison of measures such as inequality, polarization or poverty computed on surveys relatively distant in time captures, we argue, more accurately the effect of structural modifications in income distribution. Excluding cases of sudden shocks, in general these measures tend to move relatively slowly, in particular polarization. For our specific case, since we use a *relative* measure based on comparison of two distributions comparison, it becomes crucial using distributions sufficiently distant in time in order to see significant differences.

Comparisons over time, however, can be made difficult or even impossible by changes in data collection methodology (Tarozzi, 2007). In particular for what concerns survey data, there is increasing empirical evidence that questionnaires revisions can affect respondents' response in relevant ways (*inter alia* Deaton and Grosh, 2000). For example, the choice of recall period (7, 30 or x days before the interview) or the disaggregation of the expenditure items can deeply influence reports on expenditure. Other changes such as the switch from a diary-based collection to a recall-based collection can dramatically change aggregate food consumption expenditures, a relevant component of total expenditures in many developing countries.

Beegle *et al.* (2010), for example, find that in Tanzania recall modules measure lower consumption than a personal diary, with larger gaps among poorer households and for households

with more adult members. [Ahmed *et al.* \(2014\)](#), looking at Bangladesh data, also find that a switch from diary to recall reduces consumption aggregates simply because households remember their expenses better when entering them regularly in a diary. Therefore, switching the data collection methods from diary to recall likely makes poverty estimates incomparable with those of previous rounds in which consumption data were collected by diary.

In Nigeria, The Nigeria National Bureau of Statistics (NBS) uses National Living Standard Survey (NLSS) 2003/2004 and Harmonized National Living Standard Survey (HNLSS) 2009/2010 to monitor progress in poverty reduction in the country. These surveys are representative at state level, use a month-long diary to collect consumption and enumerators were in the field over a period ranging from October to September of the following year. NBS also conducts other household surveys, most notably the General Household Survey (GHS) cross section and panel.

The GHS cross section is a survey of 22,000 households carried out periodically throughout the country. It is freely downloadable from the NBS's website upon request. Available datasets include 6 rounds, from 2004/2005 to 2010/2011. Enumerators visit households once, generally in March, and ask a very standard set of questions. Data on consumption are collected by asking the household about broad categories of consumed items in the last month: food, healthcare, school, and so forth. In 2004/2005 and 2010/2011 data on consumption were not collected.

The GHS panel is a randomly selected sub-sample from GHS cross section consisting of 5,000 households. The panel covers the period 2010/2011 (Wave 1) and 2012/2013 (Wave 2). It is representative at national and zonal (geo-political) levels. Besides the questions asked in a normal GHS survey, it contains data on agricultural activities and other household income activities. Consumption data are collected using a 7-day recall period. In every panel wave, households are interviewed two times: once in the “post-planting” period, ranging from August to November, and once in the “post-harvesting” period, ranging from February to April.

Consumption data—the welfare measure we use for our analysis—from these three different sources are not directly comparable. Preliminary results based on poverty and inequality figures computed on the GHS panel and the HNLSS indicate that the figures computed using the former look substantially different from those computed on the latter. The need for comparable data requires, thus, some form of homogenization of consumption figures. In a preliminary version of this paper we focused on the 2010–2013 panel dataset. The main caveat was that a two-year difference was a short period to detect substantial modifications in the distribution. In

particular, while consumption polarization might vary due to a number of exogenous factors (crisis, shocks, etc.) and we might observe significant differences, more difficult is linking these transformations to specific covariates such as education, labor market access, spatial divide, and so forth.

In order to enable the data comparison over a longer time span (a decade), we employ *survey-to-survey* imputation techniques. Specifically, we use Wave 1 of panel data to impute consumption on the 2003/2004 NLSS survey. Given the importance of obtaining accurate estimates that are comparable over time, it is crucial to calibrate models in a year when both household consumption data and non-consumption data are available, and then use the model to impute household consumption data for years when only the non-consumption data are available. As we will discuss in greater detail below, we will use the panel Wave 2 as a benchmark to check the accuracy of our prediction and in a second stage use the same model to impute the 2003/2004 data.

The imputation process is a simplified version of the methodology developed in [Elbers *et al.* \(2003\)](#). [Stifel and Christiaensen \(2007\)](#) provide theoretical guidance regarding variables to be included in imputation models. They recommend including covariates that change over time, but call for excluding variables whose rates of return are likely to change markedly in the face of evolving economic conditions. Following [Stifel and Christiaensen \(2007\)](#), we included several household durables but excluded mobile phones, as their relationship with total household expenditure has been changing rapidly in the last ten years. In fact, ten years ago ownership of mobile phones was a good predictor of high income; today, such phones are prevalent among the lower- and middle-income classes and even among the poor ([Ahmed *et al.*, 2014](#)). Other variables include household characteristics and location; we excluded education levels since over ten years the returns on this variable might have changed too fast.

The procedure follows two stages. First, we estimate a model of log per capita expenditures on a sample from panel Wave 1. The model can be defined as:

$$\ln Y_{ik} = \alpha + \beta X_{ik} + \gamma Z_{ik} + (\eta_k + \epsilon_{ik}), \quad (1)$$

where α is an intercept, X_{ik} is the vector of explanatory variables for household i and location k , β is the vector of regression coefficients, Z is the vector of location specific variables, γ is the vector of coefficients, and the residual is decomposed into two independent components:

the cluster-specific effect, η_k , and a household-specific effect, ϵ_{ik} . This structure allows both a location effect—common to all households in the same area—and heteroskedasticity in the household-specific errors.

Second, to control for this location effect and heteroskedasticity we draw errors from the distribution of residuals for households with similar assets and in the same zone. We divide the sample into 10 groups based on deciles of a wealth index (Ferreira *et al.*, 2011) and six macro zones. The sample of the target distribution is also divided in 60 groups by the same methodology used for the original sample. Residuals are then drawn and imputed to household within each of the sixty groups. Following the bootstrap principle, residuals distribution is drawn for a number $R = 200$ of replications so as to obtain a number R of distributions. In order to preserve the statistical features of the distribution, such as the standard deviation, we selected among the 200 replications the distribution having the median standard deviation.

We apply different procedures to test the validity of the model: first, by means of in-sample criteria, *i.e.* by evaluating the R^2 size of the predicting model (1); then using out-of-sample ones, by testing the predictive capacity of the model on a known consumption distribution (2012/2013) by quantile-to-quantile analysis and other visually oriented techniques such as kernel density comparison.

Figures 1 and 2 compare results using various imputation methods. The model in Equation (1) using the above-mentioned error sampling and imputation method is compared to two alternative methods in its ability to simulate the 2012/2013 consumption distribution: the Gaussian normal regression imputation method (MI_REG)¹ and the predictive mean matching imputation method (MI_PMM)². In Figure 1 the three methods are compared via the quantile-to-quantile plot.

[Figure 1 about here.]

Our method (labelled POV_MAP) better minimizes the distance between real 2012/2013 distribution and the simulated one, in particular for values located in the upper tail of the distribution. As an additional robustness test, in Figure 2 we compare the kernel density kernel of the 2012/2013 consumption distribution (ORIG), the poverty map simulation and the the two multiple imputation outcomes.

[Figure 2 about here.]

¹<http://www.stata.com/manuals13/mimiimputeregress.pdf>.

²<http://www.stata.com/manuals13/mimiimputepmm.pdf>.

As it emerges from the graph, all the simulations are quite accurate in predicting the original data but, differently from MI outcomes, our model doesn't truncate the distribution and reproduces almost perfectly the upper tail.

3 Measuring Distributional Polarization in Nigeria: The Method Based on the Relative Distribution

3.1 Some Background on the Income Polarization Literature

Over the last two decades, the issue of polarization has come to be assigned increasing importance in the analysis of income distribution. Notwithstanding the pains the polarization literature has suffered to distinguish itself from pure inequality measurement—see *e.g.* Foster and Wolfson (1992), Levy and Murnane (1992), Esteban and Ray (1994) and Wolfson (1994, 1997), it now seems to be fairly widely accepted that polarization is a distinct concept from inequality.

Broadly speaking, the notion of polarization is concerned with the disappearance of the middle class, which occurs when there is a tendency to concentrate in the tails—rather than the middle—of the income distribution. One of the main reasons for looking at income polarization this way, which is usually referred to as “bi-polarization”, is that a well-off middle class is important to every society because it contributes significantly to economic growth, as well as to social and political stability (*e.g.* Easterly, 2001, and Pressman, 2007). In contrast, a society with high degree of income polarization may give rise to social conflicts and tensions. Therefore, in order for such risks to be minimized, it is necessary to monitor the economic evolution of the society using indices that look at the dispersion of the income distribution from the middle toward either or both of the two tails³. Measures of income polarization that correspond to this case have been proposed in the literature by Foster and Wolfson (1992), Wolfson (1994,

³More precisely, there are two characteristics that are considered as being intrinsic to the notion of bi-polarization. The first one, “increased spread”, implies that moving from the central value (median) to the extreme points of the income distribution makes the distribution more polarized than before. In other words, increments (reductions) in incomes above (below) the median will widen the distribution, that is extend the distance between the groups below and above the median and hence increase the degree of bi-polarization. On the other hand, “increased bi-polarity” refers to the case where incomes on the same side of the median get closer to each other. Since the distance between the incomes below or above the median has been reduced, this is assumed to increase bi-polarization. Thus, bi-polarization involves both an inequality-like component, the “increased spread” principle, which increases both inequality and polarization, and an equality-like component, the “increased bi-polarity” criterion, which increases polarization but lowers any inequality measure that fulfills the Pigou-Dalton transfer principle—the requirement under which inequality decreases when a transfer is made from a richer to a poorer individual without reversing their pairwise ranking. This shows that although there is complementarity between polarization and inequality, there are differences as well. See the references cited in the main text for a thorough discussion.

1997), Wang and Tsui (2000), Chakravarty and Majumder (2001), Rodríguez and Salas (2003), Chakravarty *et al.* (2007), Silber *et al.* (2007), Chakravarty (2009), Chakravarty and D'Ambrosio (2010), Lasso de la Vega *et al.* (2010), and others.

A more general notion of income polarization, which was originally proposed by Esteban and Ray (1994), regards the latter as “clustering” of a population around two or more poles of the distribution, irrespective of where they are located along the income scale. The notion of income polarization in a multi-group context is an attempt at capturing the degree of potential conflict inherent in a given distribution (see Esteban and Ray, 1999, 2008, 2011). The idea is to consider society as an amalgamation of groups, where the individuals in a group share similar attributes with the other members (*i.e.* have a mutual sense of “identification”) but in terms of the same attributes they are different from the members of the other groups (*i.e.* have a feeling of “alienation”). Political or social conflict is therefore more likely the more homogenous and separate the groups are, that is when the within-group income distribution is more clustered around its local mean and the between-group income distance is longer. In addition to Esteban and Ray (1994), indices regarding the concept of income polarization as conflict among groups have been investigated, among others, by Gradín (2000), Milanovic (2000), D'Ambrosio (2001), Zhang and Kanbur (2001), Reynal-Querol (2002), Duclos *et al.* (2004), Lasso de la Vega and Urrutia (2006), Esteban *et al.* (2007), Gigliarano and Mosler (2009) and Poggi and Silber (2010).

Much of the literature so far considered has analyzed summary measures of income polarization. Another strand uses kernel density estimation and mixture models in order to describe changes in polarization patterns over time, not just of personal incomes (as in Jenkins, 1995, 1996, Pittau and Zelli, 2001, 2004, 2006, and Conti *et al.*, 2006) but also of the cross-country distribution of per capita income (see Quah, 1996a,b, 1997, Bianchi, 1997, Jones, 1997, Paap and van Dijk, 1998, Johnson, 2000, Holzmann *et al.*, 2007, Henderson *et al.*, 2008, Pittau *et al.*, 2010, Anderson *et al.*, 2012, and others). The analysis of the shape of the income distribution provides indeed a picture from which at least three important distributional features can be observed simultaneously (Cowell *et al.*, 1996): income levels and changes in the location of the distribution as a whole; income inequality and changes in the spread of the distribution; clumping and polarization as well as changes in patterns of clustering at different modes. Finally, a rather recent (yet non-parametric) approach that combines the strengths of summary polarization indices with the details of distributional change offered by the kernel density estimates—the so-called “relative distribution”—has been employed by Alderson *et al.* (2005), Massari (2009),

Massari *et al.* (2009a,b), Borraz *et al.* (2011) and Alderson and Doran (2011, 2013) to assess the evolution of the middle class and the degree of household income polarization in a number of middle- and high-income countries in the world.

3.2 The Relative Distribution: Basic Concepts

In the current application, the *relative distribution* approach has some important advantages over the other mentioned methods of investigating income polarization. First, it readily lends itself to simple and informative graphical displays of relative data that reveal precisely where and by how much an income distribution changed over time. Second, by providing the potential for decomposition into location and shape components, it allows one to examine several hypotheses regarding the origins of distributional change—such as whether the change consists of an equal absolute subtraction or addition to all incomes that moves the overall distribution either to the left or to the right (while leaving the shape unaltered) or of shape modifications which, by definition, are independent of location shifts⁴. Lastly, it allows us to quantify the degree of polarization due to changes in distributional shape only (*i.e.* net of location shifts), thus enabling one to isolate aspects of inter-distributional inequality that are often hidden when also changes in location are examined.

Basically, the relative distribution method can be applied whenever the distribution of some quantity across two populations is to be compared, either cross-sectionally or over time⁵. To proceed, it is necessary to single out one of the two populations, refer to it as the “comparison” population, and refer to the other as the “reference” population. More formally, let Y_0 be the income variable for the reference population and Y the income variable for the comparison population. The relative distribution of Y to Y_0 is defined as the distribution of the random variable:

$$R = F_0(Y), \quad (2)$$

which is obtained from Y by transforming it by the cumulative distribution function of Y_0 , F_0 . As a random variable, R is continuous on the outcome space $[0, 1]$, and its realizations, r , are

⁴Of course, both the location and shape effects—named respectively as “growth” and “inequality” (or “distributional”) effect (Kakwani, 1993; Bourguignon, 2003, 2004)—may also concur together in producing the distributional change.

⁵Here we limit ourselves to illustrating the basic concepts behind the use of the relative distribution method. Interested readers are referred to Hancock and Morris (1998, 1999; but see also Hao and Naiman, 2010, ch. 5) for a more detailed explication and a discussion of the relationship to alternative econometric methods for measuring distributional differences. A method very similar in spirit to the relative distribution has recently been developed by Silber and Deutsch (2012).

referred to as “relative data”. Intuitively, the relative data can be interpreted as the set of positions that the income observations of the comparison population would have if they were located in the income distribution of the reference population. The probability density function of R , which is called the “relative density”, can be obtained as the ratio of the density of the comparison population to the density of the reference population evaluated at the relative data r :

$$g(r) = \frac{f(F_0^{-1}(r))}{f_0(F_0^{-1}(r))} = \frac{f(y_r)}{f_0(y_r)}, \quad 0 \leq r \leq 1, \quad y_r \geq 0, \quad (3)$$

where $f(\cdot)$ and $f_0(\cdot)$ denote the density functions of Y and Y_0 , respectively, and $y_r = F_0^{-1}(r)$ is the quantile function of Y_0 . The relative density has a simple interpretation, as it describes where households at various quantiles in the comparison distribution are concentrated in terms of the quantiles of the reference distribution. As for any density function, it integrates to 1 over the unit interval, and the area under the curve between two values r_1 and r_2 is the proportion of the comparison population whose income values lie between the r_1^{th} and r_2^{th} quantiles of the reference population.

When the relative density function shows values near to 1, it means that the two populations have a similar density at the r^{th} quantile of the reference population, and thus R has a uniform distribution in the interval $[0, 1]$. A relative density greater than 1 means that the comparison population has more density than the reference population at the r^{th} quantile of the latter. Finally, a relative density function less than 1 indicates the opposite. In this way one can distinguish between growth, stability or decline at specific points of the income distribution.

3.3 The Location/Shape Decomposition of the Relative Distribution

As we have said before, one of the major advantages of this method is the possibility to decompose the relative distribution into changes in location, usually associated with changes in the median (or mean) of the income distribution, and changes in shape (including differences in variance, asymmetry and/or other distributional characteristics) that could be linked with several factors like, for instance, polarization. Formally, the decomposition can be written as:

$$g(r) = \underbrace{\frac{f(y_r)}{f_0(y_r)}}_{\text{Overall relative density}} = \underbrace{\frac{f_{0L}(y_r)}{f_0(y_r)}}_{\text{Density ratio for the location effect}} \times \underbrace{\frac{f(y_r)}{f_{0L}(y_r)}}_{\text{Density ratio for the shape effect}}, \quad (4)$$

where $f_{0L}(y_r) = f_0(y_r + \rho)$ is a density function adjusted by an additive shift with the same shape as the reference distribution but with the median of the comparison one⁶. The value ρ is the difference between the medians of the comparison and reference distributions. If the latter two distributions have the same median, the density ratio for location differences is uniform in $[0, 1]$. Conversely, if the two distributions have different median, the “location effect” is increasing (decreasing) in r if the comparison median is higher (lower) than the reference one. The second term, which is the “shape effect”, represents the relative density net of the location effect and is useful to isolate movements (re-distribution) occurring between the reference and comparison populations. For instance, we could observe a shape effect function with some sort of (inverse) U-shaped pattern if the comparison distribution is relatively (less) more spread around the median than the location-adjusted one. Thus, it is possible to determine whether there is polarization of the income distribution (increases in both tails), “downgrading” (increases in the lower tail), “upgrading” (increases in the upper tail) or convergence of incomes towards the median (decreases in both tails).

3.4 Relative Polarization Indices

The relative distribution approach also includes a *median relative polarization* index (MRP), which is based on changes in the shape of the income distribution to account for polarization. This index is normalized so that it varies between -1 and 1, with 0 representing no change in the income distribution relative to the reference year. Positive values represent more polarization—*i.e.* increases in the tails of the distribution—and negative values represent less polarization—*i.e.* convergence towards the center of the distribution. The MRP index for the comparison population can be estimated as (Morris *et al.*, 1994, p. 217):

$$\text{MRP} = \frac{4}{n} \left(\sum_{i=1}^n \left| r_i - \frac{1}{2} \right| \right) - 1, \quad (5)$$

where r_i is the proportion of the median-adjusted reference incomes that are less than the i^{th} income from the comparison sample, for $i = 1, \dots, n$, and n is the sample size of the comparison population.

⁶Median adjustment is preferred here to mean adjustment because of the well-known drawbacks of the mean when distributions are skewed. A *multiplicative* median shift can also be applied. However, the multiplicative shift has the drawback of affecting the shape of the distribution. Indeed, the equi-proportionate income changes increase the variance and the rightward shift of the distribution is accompanied by a flattening (or shrinking) of its shape (see *e.g.* Jenkins and Van Kerm, 2005).

The MRP index can be additively decomposed into the contributions to overall polarization made by the lower and upper halves of the median-adjusted relative distribution, enabling one to distinguish downgrading from upgrading. In terms of data, the *lower relative polarization* index (LRP) and the *upper relative polarization* index (URP) can be calculated as follows:

$$\text{LRP} = \frac{8}{n} \left[\sum_{i=1}^{n/2} \left(\frac{1}{2} - r_i \right) \right] - 1, \quad (6)$$

$$\text{URP} = \frac{8}{n} \left[\sum_{i=n/2+1}^n \left(r_i - \frac{1}{2} \right) \right] - 1, \quad (7)$$

with $\text{MRP} = \frac{1}{2} (\text{LRP} + \text{URP})$. As the MRP, LRP and URP range from -1 to 1, and equal 0 when there is no change.

3.5 Adjustment for Covariates

Similarly to what is observed for location and shape decomposition, it is possible to adjust the relative distribution for changes in the distribution of covariates measured on the households, which often vary systematically by population. The covariate adjustment technique can be used to separate the impacts of changes in population composition from changes in the covariate-response relationship⁷. This decomposition according to covariates draws on the definition of a counter-factual distribution for the response variable in the reference population that is *composition-adjusted* to have the same distribution of the covariates as the comparison population.

Assume for simplicity that the covariate Z is categorical⁸. Let $\{\pi_k^0\}_{k=1}^K$ and $\{\pi_k\}_{k=1}^K$, where K is the number of categories of the covariate, denote the probability mass functions of Z for the reference and comparison populations, *i.e.* their composition according to the covariate. For conditional comparisons of the response variable Y across the two populations one can consider

⁷Recently, there have been several papers that have studied decomposition methods to explain changes in the unconditional distribution of an outcome variable due to either changes in the distribution of the covariates, or changes in the conditional distribution of the outcome given covariates, or both—see for instance the extensive survey by Fortin *et al.* (2011) on the wage decomposition literature. Benefits and drawbacks of some of these methods, and how they are often largely subsumed by the relative distribution framework, are reviewed in Handcock and Morris (1999, ch. 2).

⁸The extensions to continuous and multivariate covariates are considered in Handcock and Morris (1999, ch. 7).

the density of Y_0 given that $Z_0 = k$:

$$f_{Y_0|Z_0}(y|k), \quad k = 1, \dots, K, \quad (8)$$

and the density of Y given that $Z = k$:

$$f_{Y|Z}(y|k), \quad k = 1, \dots, K. \quad (9)$$

These densities represent the covariate-response relationship. The marginal densities of Y_0 and Y can be written, respectively, as:

$$f_0(y) = \sum_{k=1}^K \pi_k^0 f_{Y_0|Z_0}(y|k) \quad \text{and} \quad f(y) = \sum_{k=1}^K \pi_k f_{Y|Z}(y|k). \quad (10)$$

Then, the counter-factual distribution with the covariate composition of the comparison population and the covariate-response relationship of the reference population is:

$$f_{0C}(y) = \sum_{k=1}^K \pi_k f_{Y_0|Z_0}(y|k), \quad (11)$$

and can be used to decompose the overall relative distribution into a component that represents the effect of changes in the marginal distribution of the covariate (the “composition effect”) and a component that represents the changes in the covariate-response relationship (the “residual effect”). The decomposition can be represented in the following terms:

$$g(r) = \underbrace{\frac{f(y_r)}{f_0(y_r)}}_{\text{Overall relative density}} = \underbrace{\frac{f_{0C}(y_r)}{f_0(y_r)}}_{\text{Density ratio for the composition effect}} \times \underbrace{\frac{f(y_r)}{f_{0C}(y_r)}}_{\text{Density ratio for the residual effect}}. \quad (12)$$

Comparison of $f(y_r)$ to $f_{0C}(y_r)$ —*i.e.* the residual effect—holds the population composition constant, and therefore isolate changes of income distribution due to the fact that returns to the selected covariate changed over time. By contrast, $f_{0C}(y_r)$ and $f_0(y_r)$ have the same covariate-response relationship, and the comparison between them—*i.e.* the composition effect—*isolate* the changes due to the different composition of the population under the assumption that the conditional distribution of income remain unchanged.

4 Results

4.1 Changes in the Nigerian Household Consumption Distribution

Table 1 provides summary measures for household total consumption expenditure per capita in 2003/2004 and 2012/2013.

[Table 1 about here.]

Besides the growth of the real mean and median consumption expenditures, the most notable feature is that consumption shares of the poorest percentiles of the population decreased between approximately 0.4 and 0.6 percent a year in the period examined, in contrast to what is observed for the richest percentiles, whose shares experienced average yearly increases of around 2 percent. The Gini index grew at an annual average rate of 1.3 percent between 2003/2004 and 2012/2013, while the increment in inequality detected by the Theil index is more pronounced, with an average growth rate of 5.3 percent per annum. As for polarization, a slight increase is detected by both the Foster-Wolfson (1992) and Duclos-Esteban-Ray (2004) measures, which amounts to around 0.3 percent per year in the first case and almost 0.9 percent in the second⁹.

Further insight on the key changes occurring in the distribution of total per capita consumption expenditure of Nigerian households is provided by Figure 3(a), which shows the density overlay for the two survey waves¹⁰.

[Figure 3 about here.]

Two major observations are apparent from this figure: first, the whole distribution shifted rightward following the increment in the median, and second, there was also an alteration of the shape—the consumption distribution is in fact more dispersed in 2012/2013 than in 2003/2004, as it appears to be characterized by a smaller peak and a fatter upper tail that are quite visible in the density overlay. The declines in the mass at the lower and middle ranges of the distribution, and the concomitant spreading out of expenditures in its top half, are also noticeable from Table 1, where the reported values of the standard deviation, skewness, and kurtosis all show a remarkable growth from one survey wave to the next.

⁹The Foster-Wolfson and the Duclos-Esteban-Ray polarization measures have been estimated using the latest version of DASP, the Distributive Analysis Stata Package (Araar and Duclos, 2013), which is freely available at <http://dasp.ecn.ulaval.ca/>.

¹⁰To handle data sparseness, the two densities have been obtained by using an adaptive kernel estimator with a Silverman's plug-in estimate for the pilot bandwidth (see *e.g.* Van Kerm, 2003). The advantage of this estimator is that it does not over-smooth the distribution in zones of high expenditure concentration, while keeping the variability of the estimates low where data are scarce—as, for example, in the highest expenditure ranges.

However, the graphical display above does not provide much information on the relative impact that location and shape changes had on the differences in the two distributions at every point of the expenditure scale. It also does not convey whether the upper and lower tails of the consumption distribution were growing at the same rate and for what reasons (*i.e.*, location and/or shape driven). As already pointed out in Section 3, this is exactly what the relative distribution method is particularly good at pulling out of the data.

The full distribution of total per capita consumption expenditure of Nigerian households in 2012/2013 relative to the 2003/2004 reference distribution is examined in Figure 3(b). This plot shows the fraction of households in 2012/2013 that fall into each percentile of the 2003/2004 consumption distribution. Changes in the distribution are indicated by the generally positive slope of the relative density¹¹, which implies a decrease of the mass of households below the 2003/2004 median expenditure. More specifically, the relative distribution is less than 1 for $r \leq 0.57$ and more than 1 for $r > 0.57$. This means that if we choose any percentile between the 1st and the 57th in the 2003/2004 distribution, the fraction of households in 2012/2013 whose consumption rank corresponds to the chosen percentile is less than the analogous fraction of households in 2003/2004. Consumption growth increased instead the share of households in the top deciles of the 2012/2013 distribution, positively affecting those in deciles 7 through 10.

To get a more detailed picture, we decompose the relative density into location and shape effects according to Equation (4). Figure 3(c) presents the effect only due to the median shift, that is the pattern that the relative density would have displayed if there had been no change in distributional shape but only a location shift of the density. The effect of the median shift was quite large. This alone would have moved out of the four lowest deciles of the reference distribution a substantial fraction of 2012/2013 households and placed them in any of the remaining deciles. Note, however, that neither tail of the observed relative distribution is well reproduced by the median shift. For example, the top decile of Figure 3(c) is about 1.4, below the value of 1.9 observed in the actual data, and the bottom deciles of the same figure are also substantially lower than observed. These differences are explained by the shape effect presented in Figure 3(d), which shows the relative density net of the median influence. Without the higher median, the greater dispersion of consumption expenditures in 2012/2013 would have led to relatively more low-consuming households in 2012/2013, and this effect was mainly concentrated in the

¹¹The relative density function has been obtained by fitting a local polynomial to the estimated relative data. Throughout, we rely on the R statistical package `reldist` (Handcock, 2013) to implement the relative distribution method.

bottom decile. By contrast, at the top of the distribution the higher spread worked in the same direction of the location shift: operating by itself, it would have increased the share of households in the top decile of the 2012/2013 consumption distribution by nearly 50 percent. In sum, once changes in real median expenditure are netted out, a U-shaped relative density is observed, indicating that income (proxied by consumption) polarization was hollowing out the middle of the Nigerian household consumption distribution—with a cumulative loss that more than halved the number of households in deciles 2 through 8 of the 2012/2013 distribution.

A link between what we have observed in the graphical analysis and the quantification of the degree of polarization is captured by the relative polarization indices. These indices keep track of changes in the shape of the distribution and measure their direction and magnitude. Table 2 reports the median, lower and upper polarization indices computed from the data using Equations (5)–(7).

[Table 2 about here.]

The median index is significantly positive, implying a dispersion of the consumption distribution from the middle toward either or both of the two tails. The lower and upper polarization estimates indicate that both tails of the distribution are significantly positively polarized. The lower index, however, is slightly larger, indicating greater polarization in the lower tail of the distribution than in the upper tail.

4.2 Covariate Decompositions

So far we have focused on comparing the distribution of Nigerian household consumption expenditure between two points in time. However, there are often covariates measured on the households which vary over time, and the impact of these changes on the observed outcomes could be of interest to economic policy and suggest possibilities worthy of consideration by its designers. In the relative distribution setting, exploring the distributional impacts of changes in a covariate requires that the relative distribution is adjusted for these changes using the methods from Section 3.5. This makes it possible to separate the impacts of changes in the distribution of the covariate (the composition effect) from changes in the conditional distributions of household consumption expenditure given the covariate levels (the residual effect). Our Nigerian consumption microdata provide an opportunity to use this covariate adjustment technique as they contain a large set of covariates describing various socio-demographic characteristics of the

respondents, household assets and characteristics of the dwelling. Here, the analysis is restricted to the following covariates: sex of household head; literacy status of household head; zone; main material used for floor; main source of drinking water; main cooking fuel; main toilet facility. This selection was inspired both from previous poverty research—which advocates the inclusion of covariates that change over time, but excluding those of them that are likely to change markedly in the face of evolving economic conditions (*e.g.* [Stifel and Christiaensen, 2007](#))—and the fact that many of the covariates excluded from the analysis did not affect the statistical significance of the predicting model used to impute the 2003/2004 data.

Table 3 presents the usual summary statistics for the population subgroups defined by the levels of the covariates analyzed.

[Table 3 about here.]

The corresponding average percentage changes between 2003/2004 and 2012/2013 are given in Table 4.

[Table 4 about here.]

Both the mean and median consumption expenditures rose during the period analyzed for all population subgroups. At the same time, all of them experienced increasing inequality according to both the Gini coefficient and the Theil index. Population and income shares changed instead more heterogeneously, following patterns of increases and decreases with different magnitudes over time. In particular, there appears to have been no change in the proportions of both male-headed and female-headed households and in those of households cooking with firewood or other sources of energy. By contrast, the fractions of households with a literate head and good quality housing infrastructures (such as safe water and medium-high quality flooring material) grew considerably relative to their counterparts—households with no toilet facility, however, are more common in 2012/2013 than in 2003/2004. Finally, the proportions of households that consist of individuals living in the North East and North West zones of the country increased between 2003/2004 and 2012/2013, whereas households in the South South and the South East declined slightly; for their part, the shares of households in the North Central and South East regions are about the same in both surveys.

The above population trends are also visible from Figure 4, which plots the relative distributions of the covariates for 2012/2013 to 2003/2004.

[Figure 4 about here.]

Conceptually, these relative densities are similar to the one constructed for consumption expenditure in the previous section, though the graphs are not nearly as smooth because of natural discreteness of the covariates. By reading across the bottom axis one can see the frequencies of reference households cumulated by levels of the covariates, while reading off the y -axis for a given level of the categorical variables allows one to find the relative frequency of comparison households in each group defined by that level. The labels at the top show the categories of the covariates, and can be used for both the reference and comparison populations.

However, as already mentioned earlier in this section, in order to assess the impact of changes in population characteristics on the Nigerian consumption distribution the relative density must be decomposed by the distributions of the covariates. This is shown in Figure 5, which presents the covariate composition effects, and Figure 6, which displays the effects of residual changes—*i.e.* the expected relative density of Nigerian consumption expenditures had the covariate compositions of the 2003/2004 and 2012/2013 populations been identical.

[Figure 5 about here.]

[Figure 6 about here.]

All panels in Figure 5 are pretty close to a uniform distribution, suggesting that the observed differences in population composition according to the selected covariates had little effect on the overall changes occurred over the decade. There were slight decreases in the bottom half and tiny growth at the top of the distribution associated with some of these compositional shifts, but the observed changes have not been driven by modifications in these characteristics of the population. This perception is confirmed by the adjusted distributions graphed in Figure 6 which, in the absence of major compositional effects, are not much different than the original one depicted in Figure 3(b).

A similar conclusion can be drawn when looking at Table 5, which presents the set of relative polarization indices for each group defined by the covariates obtained by comparing their consumption distributions over time¹².

[Table 5 about here.]

If each of the group-specific polarization indices were close to 0, this would imply that after holding changes in the distributions of the covariates constant there is no residual polarization

¹²Note that by comparing the subgroup distributions over time we are effectively controlling for the compositional differences, even though no explicit composition *effect* is identified.

in consumption expenditures. The polarization we observe in the overall consumption distribution would then be due entirely to changing characteristics of the population over time. Instead, we see a different scenario. Apart from the north-eastern households, the estimates indicate a statistically significant increase of polarization in the subgroup distributions, except for households who reside in the North West region of the country and those with inadequate flooring in dwelling units, for whom some convergence toward the median is detected. In all cases, the growth of polarization stems from a shift away from the median of both tails, and this seems to happen asymmetrically, as the LRP indices are always more positive than the URPs—thus indicating more polarization in the lower than in the upper tail. Households with good flooring material in dwellings or unsafe water, instead, are more polarized in the upper than in the lower tail of their consumption distribution—or at least they are so the same way. Overall, these patterns confirm that compositional shifts contributed little to the observed consumption polarization or, in other words, holding the changes in population characteristics constant does almost nothing to reduce overall polarization¹³.

The above conclusion suggests that the main drivers of polarization are to be searched elsewhere, namely in the changes occurring over the decade in the consumption distributions of the groups defined by the covariates. While the covariate adjustment technique identifies the impact of changing population characteristics on the distribution of consumption expenditures, comparing the groups defined by the covariates directly makes it possible to analyze the changes within and between these groups' consumption distributions. As already observed, most population subgroups were both location-shifted (Tables 3 and 4) and more polarized (Table 5). To see what impact these location and shape shifts in the subgroups' distributions had on their relative positions within the overall consumption distribution, we compare the changes in deciles of the between-group relative distributions for 2003/2004 and 2012/2013 to the changes that would have occurred if only the medians or shapes of the groups had changed. More specifically, for each decile we decompose the absolute change:

$$g(C : R) - g(C_0 : R_0), \quad (13)$$

where $g(C : R)$ and $g(C_0 : R_0)$ denote respectively the relative density for comparison (C) to

¹³This finding can also serve as a check of whether the observed changes in Nigerian consumption distribution are robust to sample size variations. That is, had the modifications in population characteristics due to artifacts of the sample size, rather than to real population trends, our results would not be affected by them.

reference (R) groups of the categorical variables in 2012/2013 and 2003/2004, into the marginal effect of the median shift from the 2003/2004 relative density:

$$g(C_{0L} : R_{0L}) - g(C_0 : R_0), \quad (14)$$

and those of the shape changes in the subgroups' consumption distributions:

$$g(C_{0L} : R) - g(C_{0L} : R_{0L}), \quad (15)$$

$$g(C : R) - g(C_{0L} : R), \quad (16)$$

where R_{0L} and C_{0L} denote the distributions of the reference and comparison groups adjusted to have the same median as 2012/2013 but with the same shape as 2003/2004¹⁴. Summing up to the total difference given by Equation (13), these effects form a complete decomposition and allow us to determine what proportion of households were moved into or out of a decile of the overall distribution by changes in relative median and group-specific shape.

The spatial distribution of household consumption expenditure definitely provided the most attractive results. Figure 7 presents the decomposition for each of the six Nigerian macro-regions as compared to the rest of the country.

[Figure 7 about here.]

The solid bars show the total change by decile from Equation (13), and each of the lines represents one of the three components in the decomposition defined by Equations (14)–(16). We can see two ongoing distinctive patterns, both accentuating polarization. In the South South and the South West, relative to the rest of the country, residents tend to move out of the lower deciles of the distribution due to changes in relative median. More precisely, had the location effect been the only one operating, we would have seen in both cases a clear transition of Southerners from lower to upper deciles of the national distribution. However, the shape effect of both regions moved in the opposite direction, partially offsetting the positive impact of growth. Particularly in the lower deciles, the shape change is positive, indicating a clear trend of lower polarization in these areas that goes in the opposite direction *vis-à-vis* the national (residual) trend. This pattern is mirrored by what is going on in the upper deciles: a location effect higher than in

¹⁴The decomposition follows the spirit of that presented in [Bernhardt et al. \(1995\)](#) and [Handcock and Morris \(1998, 1999\)](#), to whom we refer the reader for more details.

the rest of the country (especially in the South West) and an accentuated tendency to upper polarization in both regions. For what concerns the North East and the North West, the conflict-stricken areas, had the location effect been the only operating force we would have seen a disproportionate increase of people in these regions occupying the lower national deciles compared to the rest of the country—they basically lagged behind. The increase of polarization in the rest of the country helped to offset this effect, filling the lower deciles of households from other regions too, whereas for the rest of the distribution we observe in practice a generalized decline of the relative position of these regions in the national distribution. Finally, while the North Central improves relative to the rest of the country in lower deciles, the South East comes to show a more articulated pattern of distributional change.

Results for the other covariates (not shown here but available upon request) looked as expected: compared to 2003/2004, households with an illiterate head or not having good cooking material, toilet, floor and safe water were all increasingly occupying the lower deciles of the distribution, and the gap in terms of consumption with the rest was increasing. Instead, the relative fraction of households headed by females in the upper deciles of the distribution was rising during the period, whereas male-headed households were moving into the deciles below the median. In spite of the fact that Nigerian society is mainly patriarchal, where men have better access to productive resources than women, the poor seems more among men than women.

5 Concluding Remarks

In the last two decades there has been two emerging narratives on Sub-Saharan Africa. The first paints a picture of an emerging continent where middle classes are expanding, and prosperity is reaching large swaths of the population ([African Development Bank, 2011](#); [Fine *et al.*, 2012](#)). The other narrative acknowledges the relatively robust growth in the past two decades, but points to slow reduction in poverty. According to this second narrative, the lack of faster reduction in poverty may be due, in part, to increasing disparities.

Nigeria, the most populous country in the African continent, experienced a stable and sustained growth over the last years, but despite this, the outcomes in terms of poverty reduction have not been satisfactory: while poverty seems to have declined in the coastal South and around the Federal Capital, Abuja, a large belt of north-eastern states have experienced a clear stagnation in poverty reduction.

Our conjecture is that Nigeria in the last decade has also gone through significant changes in the distribution of economic resources that generated mainly, but not exclusively, a fast increase of inequalities. Inequality, we argue, represents just one aspect of the whole problem: the country is also undergoing through a fast process of polarization. Polarization is increasingly becoming a concern in many developing countries. In income-polarized societies people cluster around group means and tend to be far from the mean/median of the overall distribution; as a consequence, the middle class in polarized societies struggles to consolidate its position. This has several economic consequences for the country, but also reflects in growing political instability.

Studies on polarization in Nigeria are surprisingly few and have tackled the topic with a narrow approach. This paper aims at filling this gap by undertaking an analysis that is innovative from different points of view. First, the period considered is a decade; the length of the time span, in absence of big shocks, is crucial if one wants to detect significant transformations in the welfare distribution. Second, the welfare measures compared (consumption per capita) are fully homogenous. To obtain this we made use of survey-to-survey estimation techniques (Elbers *et al.*, 2003) and tested their robustness using different methodologies and comparing results.

Finally, and most importantly, we employed the “relative distribution” approach (Handcock and Morris, 1998, 1999) to analyze changes in the Nigerian household consumption distribution in the considered period. The novelty of this method consists in providing a non-parametric framework for taking into account all of the distributional differences that could arise in the comparison of distributions over time and space. In this way, we have been able to summarize multiple features of the expenditure distribution that would not be detected easily from a comparison of standard measures of inequality and polarization.

The analysis reveals significant changes in the consumption distribution. Net of an average increase in consumption, a clear rise in polarization is detected, meaning that the distributional movements observed between 2003/2004 and 2012/2013 hollowed out the middle of the Nigerian household consumption distribution and increased concentration of the mass toward high and lowest deciles.

This pattern of distributional change, however, is not entirely homogeneous within the country, but varies from zone to zone. By means of covariates analysis, controlling by spatial characteristics of household head, we are able to highlight a relevant issue: in the South (South West and South South) relative to the rest of the country lower deciles tend to be emptied accentuating the tendency to upper polarization in both regions. In the North West and in the

conflict-stricken North East, had been the growth effect the only operating force, we would have seen a disproportionate increase of people of this region occupying the lower national deciles compared to the rest of the country.

These modifications occurred between 2003/2004 and 2012/2013 describe a situation of accentuated polarization where households living in the North increasingly moved from the center towards the bottom of consumption distribution, while Southern households increasingly moved upward. The overall impact was a generalized hollowing out of distribution center and a further accentuation of the North-South divide already characterizing the country.

Understanding the political and economic consequences of these sharp distributional changes is beyond the scope of this paper. However, some clear trends can be already foreseen. As mentioned, the non-consolidation of the middle class is the first and more obvious side effect of accentuated polarization. Second, and not so much explored in the context of developing countries, the tendency of polarized society to be more conflict-prone. Recent episodes in Brazil, Egypt and Turkey suggest the existence of this link between polarization and conflict, yet so far no relevant empirical evidence has been produced to underpin the existing theoretical models (Esteban and Ray, 1999, 2008, 2011). Nigeria is clearly an ideal candidate for such analysis, and our future research will be directed in understanding how existing conflicts in Nigerian societies can be interpreted and linked to the patterns of polarization.

References

- African Development Bank. *Annual Report*. Statistics Department of the African Development Bank Group, 2011. Available at <http://www.afdb.org/knowledge/publications/annual-report/annual-report-2011/>.
- F. C. Ahmed, C. Dorji, S. Takamatsu, and N. Yoshida. *A Hybrid Survey to Improve the Reliability of Poverty Statistics in a Cost-Effective Manner*. Mimeo, 2014.
- B. E. Aigbokhan. *Poverty, Growth and Inequality in Nigeria: A Case Study*. AERC Research Paper 102, African Economic Research Consortium, Nairobi, 2000. Available at http://www.aercafrica.org/index.php/publications/view_document/112-poverty-growth-and-inequality-in-nigeria-a-case-study.
- A. S. Alderson and K. Doran. *Global Inequality, Within-Nation Inequality, and the Changing*

- Distribution of Income in Seven Transitional and Middle-Income Societies. In C. Suter, editor, *Inequality Beyond Globalization: Economic Changes, Social Transformations, and the Dynamics of Inequality*, pages 183–200. Transaction Publishers, New Brunswick, NJ, 2011.
- A. S. Alderson and K. Doran. How Has Income Inequality Grown? The Reshaping of the Income Distribution in LIS Countries. In J. C. Gornick and M. Jäntti, editors, *Income Inequality: Economic Disparities and the Middle Class in Affluent Countries*, pages 51–74. Stanford University Press, Stanford, CA, 2013.
- A. S. Alderson, J. Beckfield, and F. Nielsen. Exactly How Has Income Inequality Changed? Patterns of Distributional Change in Core Societies. *International Journal of Comparative Sociology*, 46:405–423, 2005.
- G. Anderson, O. Linton, and T. W. Leo. A Polarization-Cohesion Perspective on Cross-Country Convergence. *Journal of Economic Growth*, 17:49–69, 2012.
- A. Araar. *On the Decomposition of Polarization Indices: Illustrations with Chinese and Nigerian Household Surveys*. Working Paper 08-06, Centre Interuniversitaire sur le Risque, les Politiques Économiques et l'Emploi (CIRPÉE), Montréal, 2008. Available at http://www.cirpee.org/fileadmin/documents/Cahiers_2008/CIRPEE08-06.pdf.
- A. Araar and J.-Y. Duclos. *User Manual for Stata Package DASP: Version 2.3*. PEP, World Bank, UNDP and Université Laval, 2013. Available at http://dasp.ecn.ulaval.ca/modules/DASP_V2.3/DASP_MANUAL_V2.3.pdf.
- T. T. Awoyemi and A. Araar. *Explaining Polarization and its Dimensions in Nigeria: A DER Decomposition Approach*. Paper prepared for the 14th Annual Conference on Econometric Modelling for Africa, Abuja, Nigeria, July 8–10, 2009. Available at http://www.africametrics.org/documents/conference09/papers/Awoyemi_Araar.pdf.
- T. T. Awoyemi, I. B. Oluwatayo, and A. Oluwakemi. *Inequality, Polarization and Poverty in Nigeria*. Working Paper 2010-04, Partnership for Economic Policy (PEP), Nairobi, 2010. Available at <http://www.pep-net.org/communications/working-papers/>.
- K. Beegle, J. De Weerdt, J. Friedman, and J. Gibson. *Methods of Household Consumption Measurement through Surveys: Experimental Results from Tanzania*. Policy Research Working

- Paper 5501, The World Bank, Washington, DC, 2010. Available at <http://go.worldbank.org/UE8VCVJXK0>.
- A. Bernhardt, M. Morris, and M. S. Handcock. Women's Gains or Mens Losses? A Closer Look at the Shrinking Gender Gap in Earnings. *American Journal of Sociology*, 101:302–328, 1995.
- M. Bianchi. Testing for Convergence: Evidence from Nonparametric Multimodality Tests. *Journal of Applied Econometrics*, 12:393–409, 1997.
- N. Birdsall. *The (Indispensable) Middle Class in Developing Countries; or, the Rich and the Rest, Not the Poor and the Rest*. Working Paper 207, Center for Global Development, Washington, DC, 2010. Available at http://www.cgdev.org/files/1423994_file_Birdsall_Indispensable_Middle_FINAL.pdf.
- F. Borraz, N. G. Pampillón, and M. Rossi. *Polarization and the Middle Class*. dECON Working Paper 20, Universidad de la República, Montevideo, 2011. Available at <http://www.fcs.edu.uy/archivos/2011.pdf>.
- F. Bourguignon. The Growth Elasticity of Poverty Reduction: Explaining Heterogeneity across Countries and Time Periods. In T. S. Eicher and S. J. Turnovsky, editors, *Inequality and Growth: Theory and Policy Implications*, pages 3–26. The MIT Press, Cambridge, MA, 2003.
- F. Bourguignon. *The Poverty-Growth-Inequality Triangle*. Working Paper 125, Indian Council for Research on International Economic Relations, New Delhi, 2004. Available at <http://www.icrier.org/pdf/wp125.pdf>.
- S. R. Chakravarty. *Inequality, Polarization and Poverty: Advances in Distributional Analysis*. Springer-Verlag Inc., New York, NY, 2009.
- S. R. Chakravarty and C. D'Ambrosio. Polarization Orderings of Income Distributions. *Review of Income and Wealth*, 56:47–64, 2010.
- S. R. Chakravarty and A. Majumder. Inequality, Polarization and Welfare: Theory and Applications. *Australian Economic Papers*, 40:1–13, 2001.
- S. R. Chakravarty, A. Majumder, and S. Roy. A Treatment of Absolute Indices of Polarization. *Japanese Economic Review*, 58:273–293, 2007.

- P. L. Conti, M. G. Pittau, and R. Zelli. Metodi non parametrici nell'analisi della distribuzione del reddito: problemi empirici ed aspetti metodologici. *Rivista di Politica Economica*, 96: 195–242, 2006.
- F. A. Cowell, S. P. Jenkins, and J. A. Litchfield. The Changing Shape of the UK Income Distribution: Kernel Density Estimates. In J. Hills, editor, *New Equalities: The Changing Distribution of Income and Wealth in the United Kingdom*, pages 49–75. Cambridge University Press, Cambridge, UK, 1996.
- C. D'Ambrosio. Household Characteristics and the Distribution of Income in Italy: An Application of Social Distance Measures. *Review of Income and Wealth*, 47:43–64, 2001.
- A. Deaton and M. Grosh. Consumption. In M. Grosh and P. Glewwe, editors, *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study*, volume 1, pages 91–133. Oxford University Press, Oxford, UK, 2000.
- J.-Y. Duclos, J.-M. Esteban, and D. Ray. Polarization: Concepts, Measurement, Estimation. *Econometrica*, 72:1737–1772, 2004.
- W. Easterly. The Middle Class Consensus and Economic Development. *Journal of Economic Growth*, 6:317–335, 2001.
- C. Elbers, J. O. Lanjouw, and P. Lanjouw. Micro-Level Estimation of Poverty and Inequality. *Econometrica*, 71:355–364, 2003.
- J.-M. Esteban and D. Ray. On the Measurement of Polarization. *Econometrica*, 62:819–851, 1994.
- J.-M. Esteban and D. Ray. Conflict and Distribution. *Journal of Economic Theory*, 87:379–415, 1999.
- J.-M. Esteban and D. Ray. Polarization, Fractionalization and Conflict. *Journal of Peace Research*, 45:163–182, 2008.
- J.-M. Esteban and D. Ray. Linking Conflict to Inequality and Polarization. *The American Economic Review*, 101:1345–1374, 2011.

- J.-M. Esteban, C. Gradín, and D. Ray. An Extension of a Measure of Polarization, with an Application to the Income Distribution of Five OECD Countries. *Journal of Economic Inequality*, 5:1–19, 2007.
- F. H. G. Ferreira, J. Gignoux, and M. Aran. Measuring Inequality of Opportunity with Imperfect Data: The Case of Turkey. *Journal of Economic Inequality*, 9:651–680, 2011.
- D. Fine, A. van Wamelen, S. Lund, A. Cabral, M. Taoufik, N. Dörr, A. Leke, C. Roxburgh, J. Schubert, and P. Cook. *Africa at Work: Job Creation and Inclusive Growth*. Report, McKinsey Global Institute, Washington, DC, 2012. Available at http://www.mckinsey.com/insights/africa/africa_at_work.
- N. Fortin, T. Lemieux, and S. Firpo. Decomposition Methods in Economics. In O. Ashenfelter and D. Card, editors, *Handbook of Labor Economics*, volume 4A, chapter 1, pages 1–102. North Holland, Amsterdam, 2011.
- J. E. Foster and M. C. Wolfson. *Polarization and the Decline of the Middle Class: Canada and the US*. OPHI Working Paper 31, University of Oxford, Oxford, 1992. Available at <http://www.ophi.org.uk/working-paper-number-31/>, now in *Journal of Economic Inequality*, 8:247–273, 2010.
- C. Gigliarano and K. Mosler. Constructing Indices of Multivariate Polarization. *Journal of Economic Inequality*, 7:435–460, 2009.
- C. Gradín. Polarization by Sub-Populations in Spain, 1973–91. *Review of Income and Wealth*, 46:457–474, 2000.
- M. S. Handcock. *Relative Distribution Methods*. Los Angeles, CA, 2013. Version 1.6-2. Project home page at <http://www.stat.ucla.edu/~handcock/RelDist>.
- M. S. Handcock and M. Morris. Relative Distribution Methods. *Sociological Methodology*, 28: 53–97, 1998.
- M. S. Handcock and M. Morris. *Relative Distribution Methods in the Social Sciences*. Springer-Verlag Inc., New York, NY, 1999.
- L. Hao and D. Q. Naiman. *Assessing Inequality*. SAGE Publications, Inc., Thousand Oaks, CA, 2010.

- D. J. Henderson, C. Parmeter, and R. R. Russell. Modes, Weighted Modes, and Calibrated Modes: Evidence of Clustering Using Modality Tests. *Journal of Applied Econometrics*, 23: 607–638, 2008.
- H. Holzmann, S. Vollmer, and J. Weisbrod. *Twin Peaks or Three Components? Analyzing the World's Cross-Country Distribution of Income*. IAI Discussion Paper 162, University of Göttingen, Göttingen, 2007. Available at <http://www2.vwl.wiso.uni-goettingen.de/ibero/papers/DB162.pdf>.
- S. P. Jenkins. Did the Middle Class Shrink During the 1980s? UK Evidence From Kernel Density Estimates. *Economic Letters*, 49:407–413, 1995.
- S. P. Jenkins. Recent Trends in the UK Income Distribution: What Happened and Why? *Oxford Review of Economic Policy*, 12:29–46, 1996.
- S. P. Jenkins and P. Van Kerm. Accounting for Income Distribution Trends: A Density Function Decomposition Approach. *Journal of Economic Inequality*, 3:43–61, 2005.
- P. A. Johnson. A Nonparametric Analysis of Income Convergence Across the US States. *Economics Letters*, 69:219–223, 2000.
- C. I. Jones. On the Evolution of the World Income Distribution. *Journal of Economic Perspectives*, 11:19–36, 1997.
- N. Kakwani. Poverty and Economic Growth with Application to Côte d'Ivoire. *Review of Income and Wealth*, 39:121–139, 1993.
- M. C. Lasso de la Vega and A. M. Urrutia. An Alternative Formulation of Esteban-Gradín-Ray Extended Measure of Polarization. *Journal of Income Distribution*, 15:42–54, 2006.
- M. C. Lasso de la Vega, A. Urrutia, and H. Díez. Unit Consistency and Bipolarization of Income Distributions. *Review of Income and Wealth*, 56:65–83, 2010.
- F. Levy and R. J. Murnane. U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations. *Journal of Economic Literature*, 30:1333–1381, 1992.
- R. Massari. *Is Income Becoming More Polarized in Italy? A Closer Look With a Distributional Approach*. DSE Working Paper 1, Sapienza University of Rome, Rome, 2009.

Available at <http://phdschool-economics.dse.uniroma1.it/website/WorkingPapers/MassariWP1.pdf>.

- R. Massari, M. G. Pittau, and R. Zelli. A Dwindling Middle Class? Italian Evidence in the 2000s. *Journal of Economic Inequality*, 7:333–350, 2009a.
- R. Massari, M. G. Pittau, and R. Zelli. Caos calmo: l'evoluzione dei redditi familiari in Italia. In L. Cappellari, P. Naticchioni, and S. Staffolani, editors, *L'Italia delle disuguaglianze*, pages 19–28. Carocci editore, Rome, 2009b.
- B. Milanovic. *A New Polarization Measure and Some Applications*. Mimeo, 2000. Available at <http://siteresources.worldbank.org/INTDECINEQ/Resources/polariz.pdf>.
- M. Morris, A. D. Bernhardt, and M. S. Handcock. Economic Inequality: New Methods for New Trends. *American Sociological Review*, 59:205–219, 1994.
- O. I. Ogunyemi. *Gender, Market Participation and Income Polarisation Among Farming and Non-Farming Households in Nigeria*. Paper presented at the World Bank Workshop on Agriculture Myths and Facts, Washington DC, US, November 20–22, 2013.
- O. I. Ogunyemi and O. A. Oni. Income Bipolarization and Poverty: Evidence from Rural Nigeria. *Journal of American Science*, 7:722–731, 2011.
- O. I. Ogunyemi, O. A. Oni, T. T. Awoyemi, and S. A. Yusuf. Income polarization and bipolarization across rural households' socio-economic features in nigeria. *World Rural Observations*, 3:62–72, 2011.
- R. Paap and H. K. van Dijk. Distribution and Mobility of Wealth of Nations. *European Economic Review*, 42:1269–1293, 1998.
- M. G. Pittau and R. Zelli. Income Distribution in Italy: A Nonparametric Analysis. *Statistical Methods and Applications*, 10:175–189, 2001.
- M. G. Pittau and R. Zelli. Testing for Changing Shapes of Income Distribution: Italian Evidence in the 1990s From Kernel Density Estimates. *Empirical Economics*, 29:415–430, 2004.
- M. G. Pittau and R. Zelli. Trends in Income Distribution in Italy: A Non-Parametric and a Semi-Parametric Analysis. *Journal of Income Distribution*, 15:90–118, 2006.

- M. G. Pittau, R. Zelli, and P. A. Johnson. Mixture Models, Convergence Clubs and Polarization. *Review of Income and Wealth*, 56:102–122, 2010.
- A. Poggi and J. Silber. On Polarization and Mobility: A Look at Polarization in the Wage-Career Profile in Italy. *Review of Income and Wealth*, 56:123–140, 2010.
- S. Pressman. The Decline of the Middle Class: An International Perspective. *Journal of Economic Issues*, 41:181–200, 2007.
- D. T. Quah. Convergence Empirics Across Economies with (Some) Capital Mobility. *Journal of Economic Growth*, 1:95–124, 1996a.
- D. T. Quah. Twin Peaks: Growth and Convergence in Models of Distribution Dynamics. *The Economic Journal*, 106:1045–1055, 1996b.
- D. T. Quah. Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs. *Journal of Economic Growth*, 2:27–59, 1997.
- M. Reynal-Querol. Ethnicity, Political Systems and Civil War. *Journal of Conflict Resolution*, 46:29–54, 2002.
- J. G. Rodríguez and R. Salas. Extended Bi-Polarization and Inequality Measures. *Research on Economic Inequality*, 9:69–83, 2003.
- J. Silber and J. Deutsch. *On Bi-Polarization and the Middle Class in Latin America: A Look at the First Decade of the Twenty First Century*. Paper presented at the 32nd General Conference of the International Association for Research in Income and Wealth, Boston, US, August 5–11, 2012. Available at <http://www.iariw.org/papers/2012/SilberPaper.pdf>.
- J. Silber, M. Hanoka, and J. Deutsch. On the Link Between the Concepts of Kurtosis and Bipolarization. *Economics Bulletin*, 4:1–6, 2007.
- D. Stifel and L. Christiaensen. Tracking Poverty Over Time in the Absence of Comparable Consumption Data. *World Bank Economic Review*, 21:317–341, 2007.
- A. Tarozzi. Calculating Comparable Statistics From Incomparable Surveys, With an Application to Poverty in India. *Journal of Business & Economic Statistics*, 25:314–336, 2007.
- P. Van Kerm. Adaptive Kernel Density Estimation. *The Stata Journal*, 3:148–156, 2003.

- Y.-Q. Wang and K.-Y. Tsui. Polarization Orderings and New Classes of Polarization Indices. *Journal of Public Economic Theory*, 2:349–363, 2000.
- M. C. Wolfson. When Inequalities Diverge. *The American Economic Review*, 84:353–358, 1994.
- M. C. Wolfson. Divergent Inequalities: Theory and Empirical Results. *Review of Income and Wealth*, 43:401–421, 1997.
- World Bank. *Nigeria. Where Has All the Growth Gone? A Poverty Update*. Report 143689-NGA, World Bank – Poverty Reduction and Economic Management Network, Washington, DC, 2013.
- X. Zhang and R. Kanbur. What Difference Do Polarisation Measures Make? An Application to China. *Journal of Development Studies*, 37:85–98, 2001.

Tables

Table 1: Summary measures of Nigerian household total consumption expenditure per capita.

	2003/2004	2012/2013
Mean	64,424	93,597
Median	49,563	64,846
Standard deviation	60,565	201,536
Skewness	7.11	53.15
Kurtosis	142.57	4,853.21
Consumption shares		
Bottom 5%	0.80	0.77
Bottom 10%	2.09	2.00
Bottom 20%	5.61	5.33
Top 20%	45.82	50.52
Top 10%	29.65	35.44
Top 5%	18.76	24.77
Inequality measures		
Gini	0.40	0.45
Theil	0.29	0.46
Polarization measures ^a		
Foster-Wolfson	0.36	0.37
Duclos-Esteban-Ray	0.24	0.26

^a The Duclos-Esteban-Ray index has been computed with the polarization sensitivity parameter α set at 0.5.

Table 2: Relative polarization indices.

Index ^a	Value	LB ^b	UB ^c	<i>p</i> -value ^d
MRP	0.18	0.16	0.19	0.00
LRP	0.21	0.18	0.24	0.00
URP	0.15	0.12	0.18	0.00

^a MRP = median relative polarization index; LRP = lower relative polarization index; URP = upper relative polarization index.

^b Lower bound of the 95 percent confidence interval.

^c Upper bound of the 95 percent confidence interval.

^d Refers to the null hypothesis of no change with respect to the reference distribution, *i.e.* that the index equals 0.

Table 3: Summary measures for Nigerian household consumption expenditure by population subgroups, 2003/2004 and 2012/2013.

	2003/2004					2012/2013						
	Mean	Median	Pop. share	Inc. share	Theil	Gini	Mean	Median	Pop. share	Inc. share	Gini	Theil
Sex of household head												
Female	75,307	61,174	0.10	0.12	0.25	0.37	114,786	85,264	0.10	0.12	0.43	0.46
Male	63,160	48,507	0.90	0.88	0.29	0.40	91,235	63,247	0.90	0.88	0.45	0.46
Literacy status of household head												
Illiterate	55,471	43,190	0.64	0.55	0.27	0.39	72,825	50,089	0.32	0.25	0.45	0.53
Literate	80,150	62,355	0.36	0.45	0.27	0.39	103,213	73,111	0.68	0.75	0.43	0.42
Zone												
North Central	45,485	34,188	0.14	0.10	0.29	0.41	87,095	61,654	0.14	0.13	0.46	0.54
North East	58,639	44,007	0.12	0.11	0.32	0.41	63,766	46,974	0.14	0.10	0.42	0.37
North West	54,018	40,185	0.25	0.21	0.29	0.40	68,026	46,628	0.26	0.19	0.44	0.70
South East	84,375	65,315	0.12	0.15	0.26	0.39	115,481	80,889	0.12	0.14	0.45	0.47
South South	66,630	49,999	0.17	0.18	0.31	0.40	108,187	79,568	0.15	0.18	0.41	0.31
South West	80,406	68,194	0.20	0.25	0.18	0.32	131,142	105,877	0.19	0.26	0.35	0.24
Main material used for floor												
Medium quality/High quality	75,162	59,539	0.61	0.71	0.26	0.38	96,899	67,985	0.91	0.94	0.44	0.45
Low quality	47,926	37,395	0.39	0.29	0.28	0.39	61,300	38,627	0.09	0.06	0.46	0.56
Main source of drinking water												
Piped/Unprotected	64,182	48,804	0.62	0.62	0.30	0.40	93,453	61,216	0.45	0.45	0.47	0.54
Protected	64,824	50,811	0.38	0.38	0.27	0.39	93,717	68,043	0.55	0.55	0.43	0.40
Main cooking fuel												
Charcoal/Kerosene/Oil/Electricity/Gas/Other	98,366	79,284	0.25	0.39	0.22	0.36	146,065	113,467	0.25	0.39	0.37	0.27
Firewood	52,927	41,990	0.75	0.61	0.26	0.38	76,126	53,770	0.75	0.61	0.43	0.51
Main toilet facility												
Flush toilet/Improved pit latrine/Uncovered pit latrine/Other	65,641	50,356	0.83	0.84	0.29	0.40	97,292	68,474	0.78	0.81	0.44	0.41
No facility	58,529	45,394	0.17	0.16	0.28	0.39	80,237	53,126	0.22	0.19	0.48	0.67

Table 4: Summary measures for Nigerian household consumption expenditure by population subgroups, average annual compound percentage changes from 2003/2004 to 2012/2013.

	Mean	Median	Pop. share	Inc. share	Gini	Theil
Sex of household head						
Female	4.79	3.76	0.00	0.00	1.68	7.01
Male	4.17	2.99	0.00	0.00	1.32	5.26
Literacy status of household head						
Illiterate	3.07	1.66	-7.41	-8.39	1.60	7.78
Literate	2.85	1.78	7.32	5.84	1.09	5.03
Zone						
North Central	7.48	6.77	0.00	2.96	1.29	7.15
North East	0.94	0.73	1.73	-1.05	0.27	1.63
North West	2.60	1.67	0.44	-1.11	1.06	10.29
South East	3.55	2.40	0.00	-0.76	1.60	6.80
South South	5.53	5.30	-1.38	0.00	0.27	0.00
South West	5.59	5.01	-0.57	0.44	1.00	3.25
Main material used for floor						
Medium quality/High quality	2.86	1.48	4.54	3.17	1.64	6.28
Low quality	2.77	0.36	-15.03	-16.06	1.85	8.01
Main source of drinking water						
Piped/Unprotected	4.26	2.55	-3.50	-3.50	1.81	6.75
Protected	4.18	3.30	4.19	4.19	1.09	4.46
Main cooking fuel						
Charcoal/Kerosene/Oil/Electricity/Gas/Other	4.49	4.06	0.00	0.00	0.30	2.30
Firewood	4.12	2.79	0.00	0.00	1.38	7.77
Main toilet facility						
Flush toilet/Improved pit latrine/Uncovered pit latrine/Other	4.47	3.47	-0.69	-0.40	1.06	3.92
No facility	3.57	1.76	2.91	1.93	2.33	10.18

Table 5: Relative polarization indices for different population subgroups.

	MRP ^a			LRP ^b			URP ^c			
	Value	LB ^d	UB ^e	Value	LB ^d	UB ^e	Value	LB ^d	UB ^e	p-value ^f
Sex of household head										
Female	0.17	0.13	0.21	0.00	0.18	0.33	0.00	0.02	0.17	0.01
Male	0.17	0.16	0.19	0.00	0.17	0.24	0.00	0.14	0.17	0.00
Literacy status of household head										
Illiterate	0.07	0.05	0.09	0.00	0.04	0.14	0.00	0.05	0.09	0.02
Literate	0.12	0.10	0.14	0.00	0.09	0.17	0.00	0.12	0.16	0.00
Zone										
North Central	0.32	0.28	0.35	0.00	0.38	0.52	0.00	0.18	0.25	0.00
North East	0.00	-0.04	0.04	0.46	-0.06	0.09	0.34	-0.01	0.06	0.37
North West	-0.05	-0.09	-0.02	0.00	-0.09	0.05	0.27	-0.08	-0.02	0.01
South East	0.17	0.13	0.20	0.00	0.15	0.30	0.00	0.10	0.18	0.00
South South	0.29	0.25	0.33	0.00	0.30	0.44	0.00	0.21	0.28	0.00
South West	0.28	0.24	0.32	0.00	0.34	0.41	0.00	0.23	0.30	0.00
Main material used for floor										
Medium quality/High quality	0.11	0.09	0.13	0.00	0.10	0.14	0.00	0.12	0.15	0.00
Low quality	-0.09	-0.13	-0.05	0.00	-0.19	-0.12	0.00	0.01	0.10	0.40
Main source of drinking water										
Piped/Unprotected	0.13	0.11	0.15	0.00	0.13	0.17	0.00	0.13	0.17	0.00
Protected	0.21	0.19	0.23	0.00	0.25	0.29	0.00	0.17	0.21	0.00
Main cooking fuel										
Charcoal/Kerosene/Oil/Electricity/Gas/Other	0.21	0.18	0.24	0.00	0.25	0.31	0.00	0.18	0.24	0.00
Firewood	0.13	0.11	0.14	0.00	0.15	0.18	0.00	0.11	0.14	0.00
Main toilet facility										
Flush toilet/Improved pit latrine/Uncovered pit latrine/Other	0.20	0.19	0.22	0.00	0.23	0.27	0.00	0.17	0.21	0.00
No facility	0.09	0.06	0.12	0.00	0.11	0.17	0.00	0.07	0.13	0.01

^a Median relative polarization index.

^b Lower relative polarization index.

^c Upper relative polarization index.

^d Lower bound of the 95 percent confidence interval.

^e Upper bound of the 95 percent confidence interval.

^f Refers to the null hypothesis of no change with respect to the reference distribution, *i.e.* that the index equals 0.

Figures

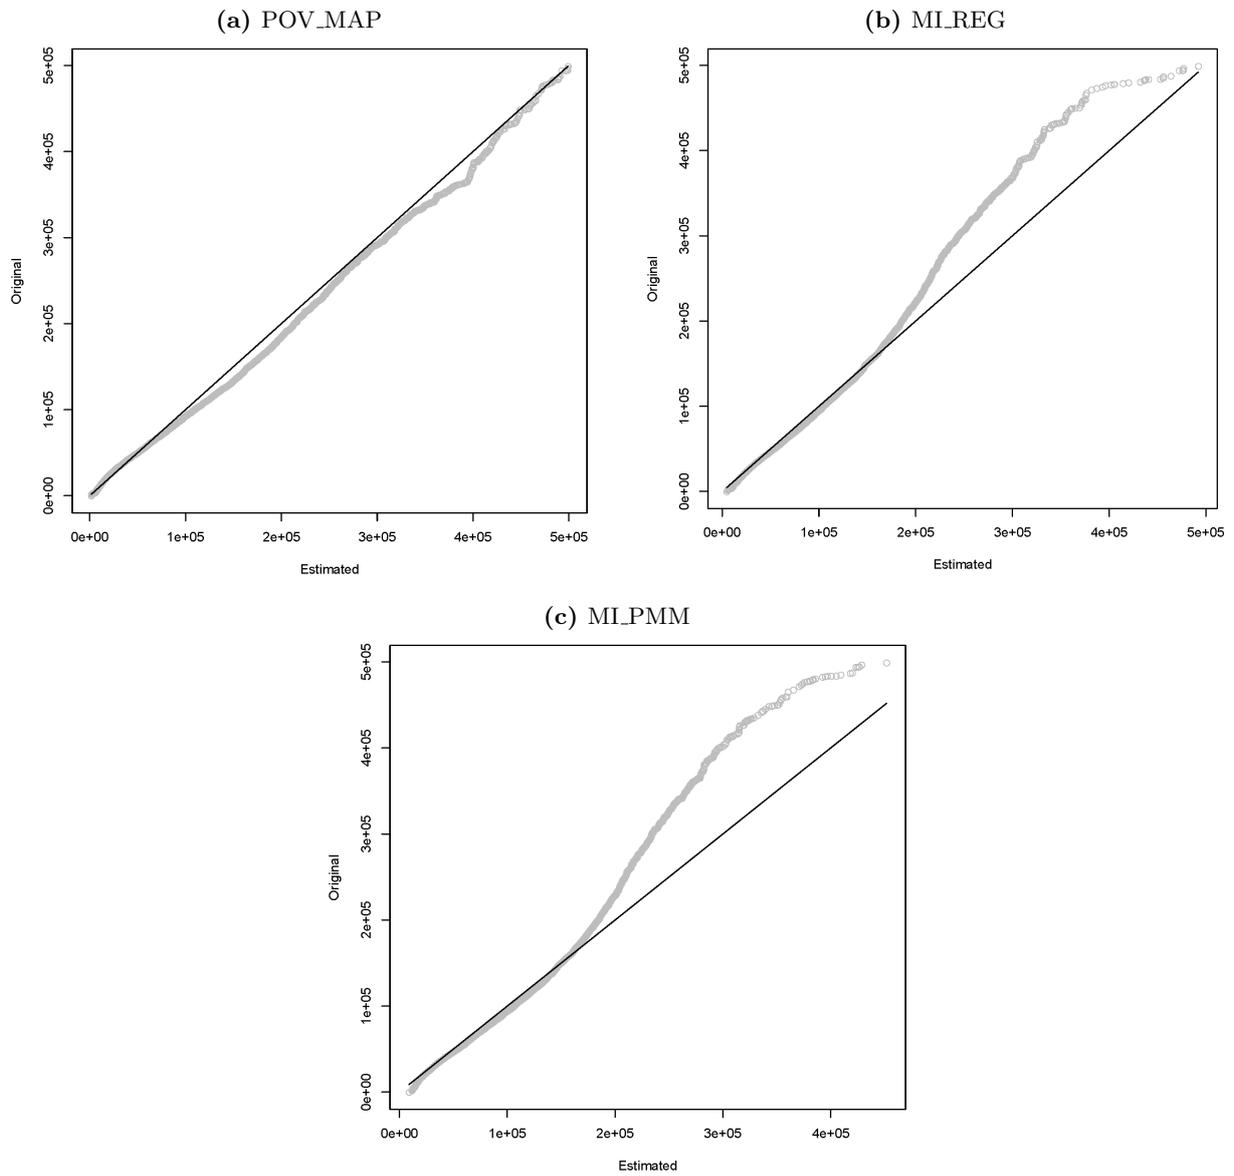
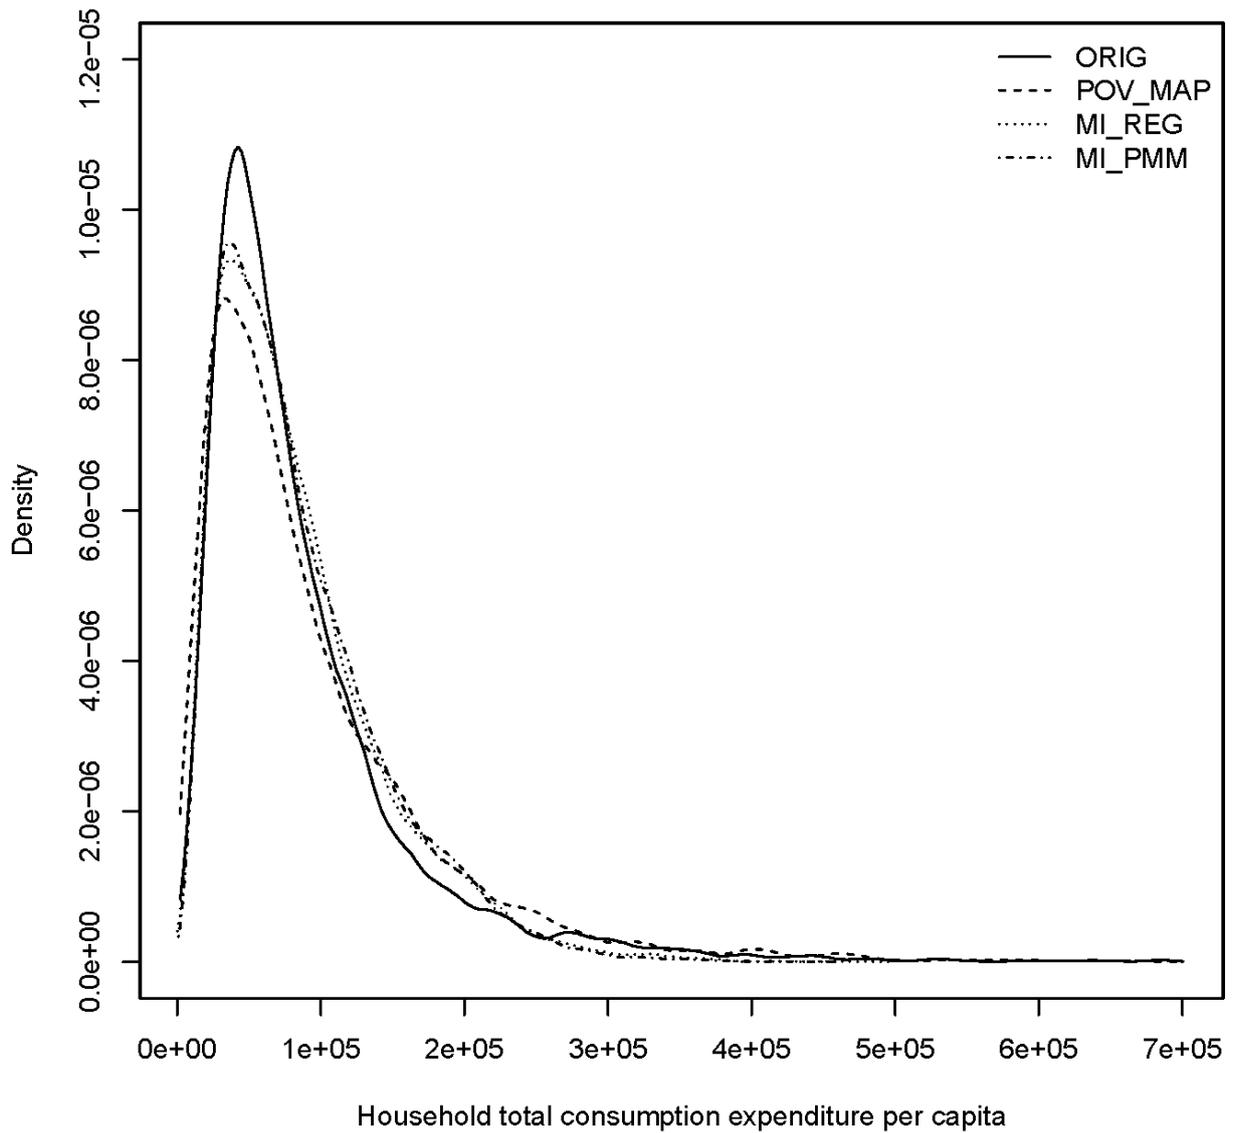


Figure 1: Quantile-to-quantile analysis.

**Figure 2:** Kernel distributions.

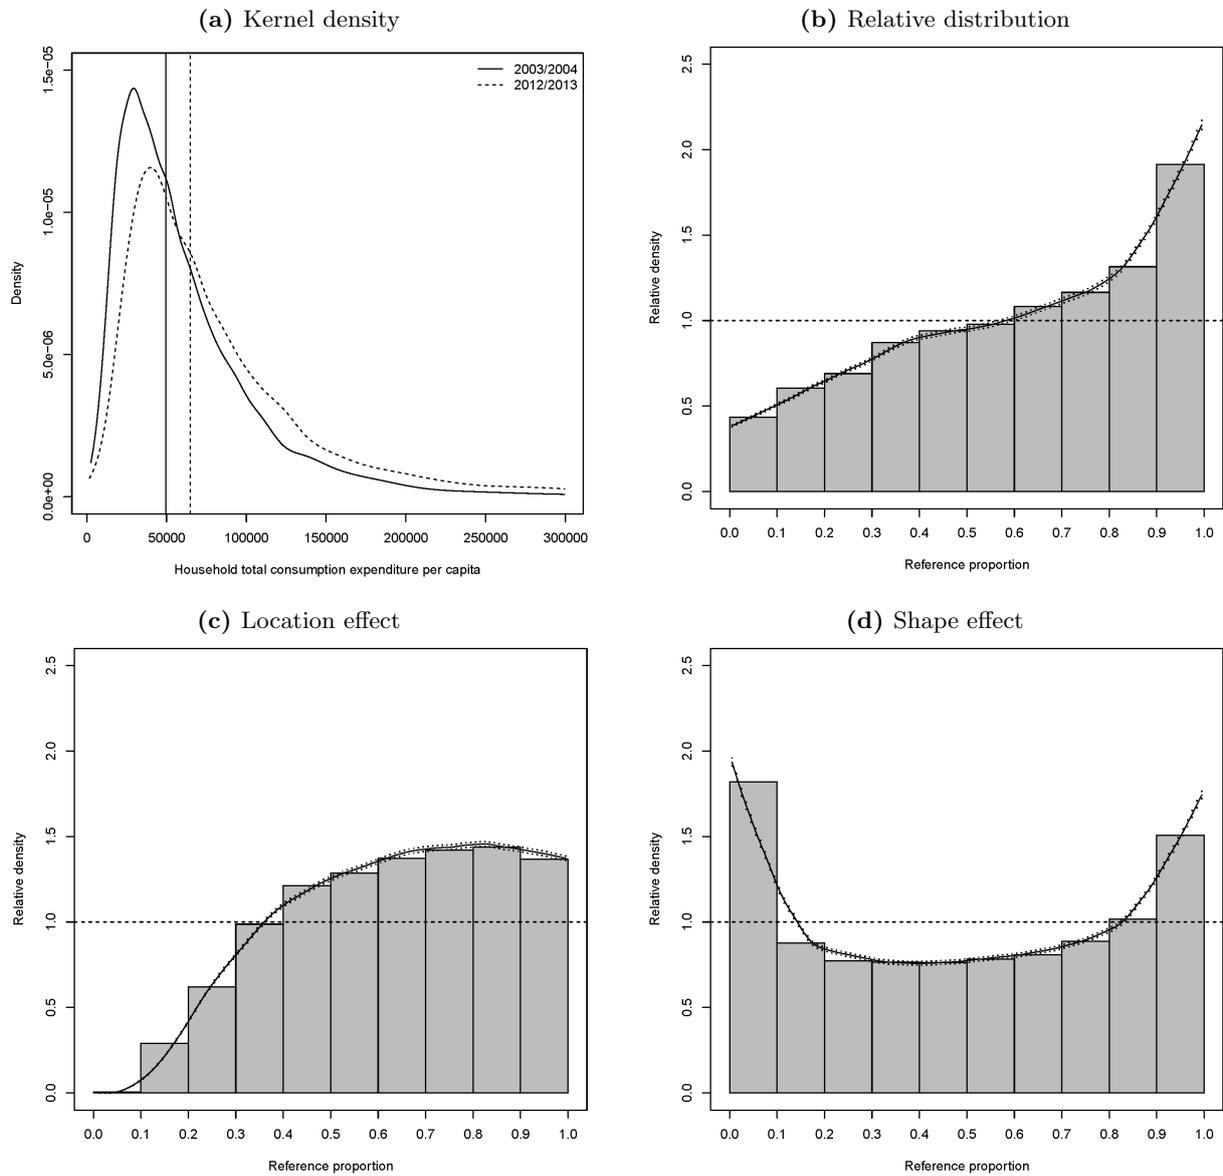


Figure 3: Changes in the Nigerian household consumption distribution between 2003/2004 and 2012/2013. In panel (a), expenditures in the upper tiers of the distributions have been truncated for better presentation of the graph, where the vertical lines denote the medians of the two survey waves. The bars in panels (b)–(d) represent the decile breakdown of the relative distribution, showing the fraction of 2012/2013 households that fall into each 2003/2004 decile, while dotted lines indicate the 95 percent pointwise confidence limits based on the asymptotic normal approximation (Handcock and Morris, 1999, p. 144).

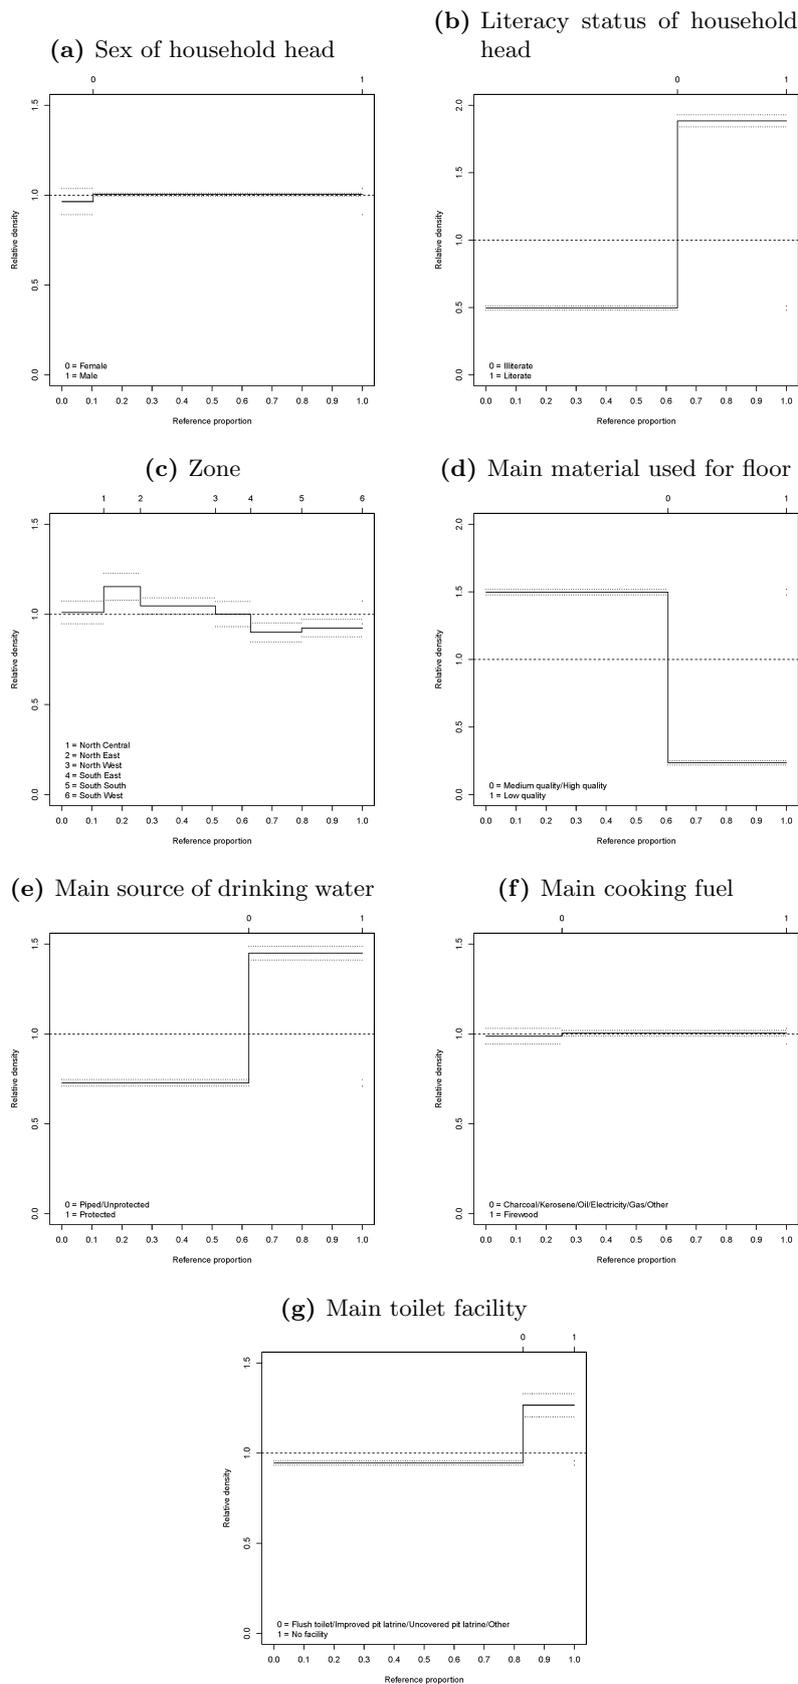


Figure 4: The relative distributions of the covariates for 2012/2013 to 2003/2004. The upper axes are labelled by the levels of categorical variables. The dotted lines are 95 percent pointwise confidence bounds.

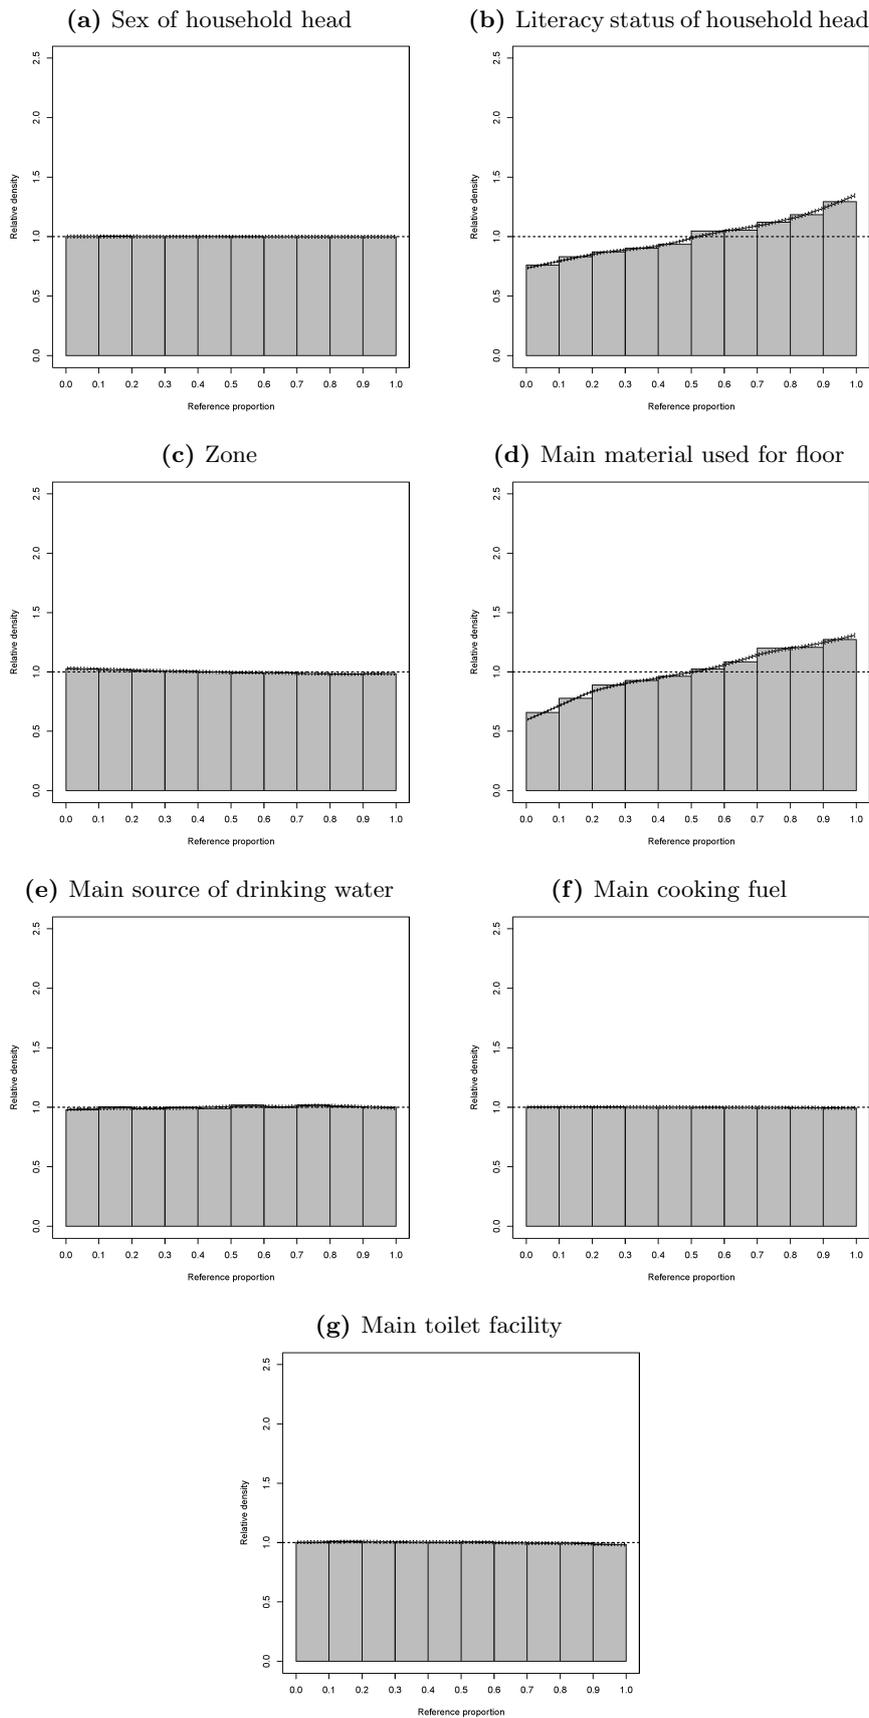


Figure 5: The effects of changes in the covariate distributions on the 2012/2013 to 2003/2004 relative density of Nigerian consumption expenditure.

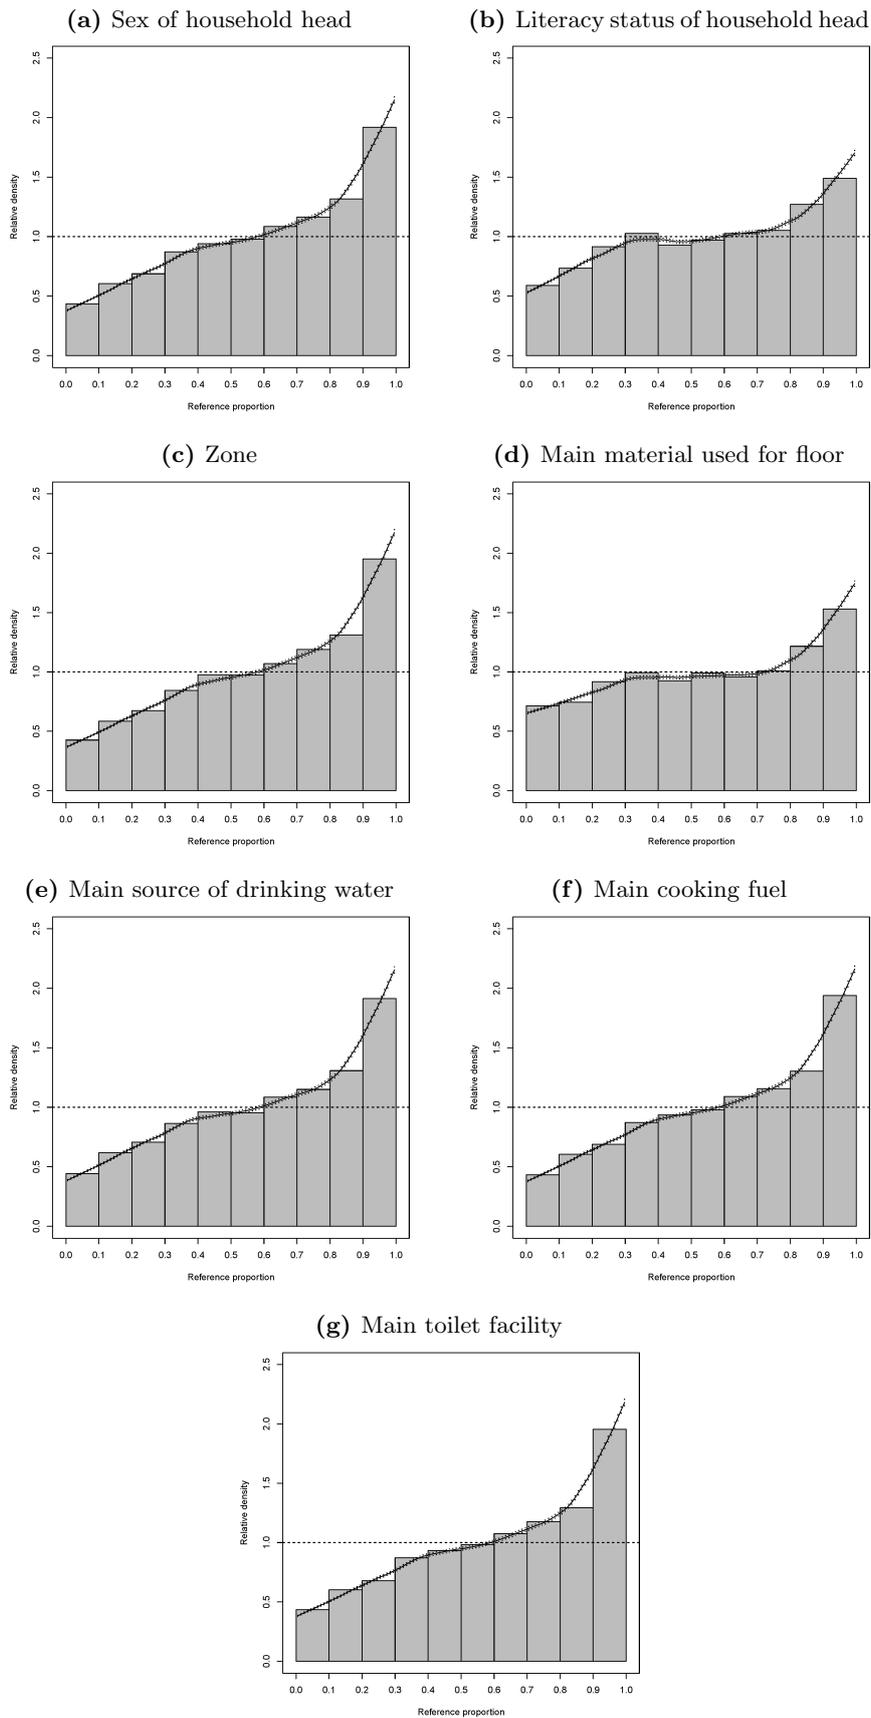


Figure 6: The composition-adjusted relative densities of Nigerian consumption expenditure for 2012/2013 to 2003/2004.

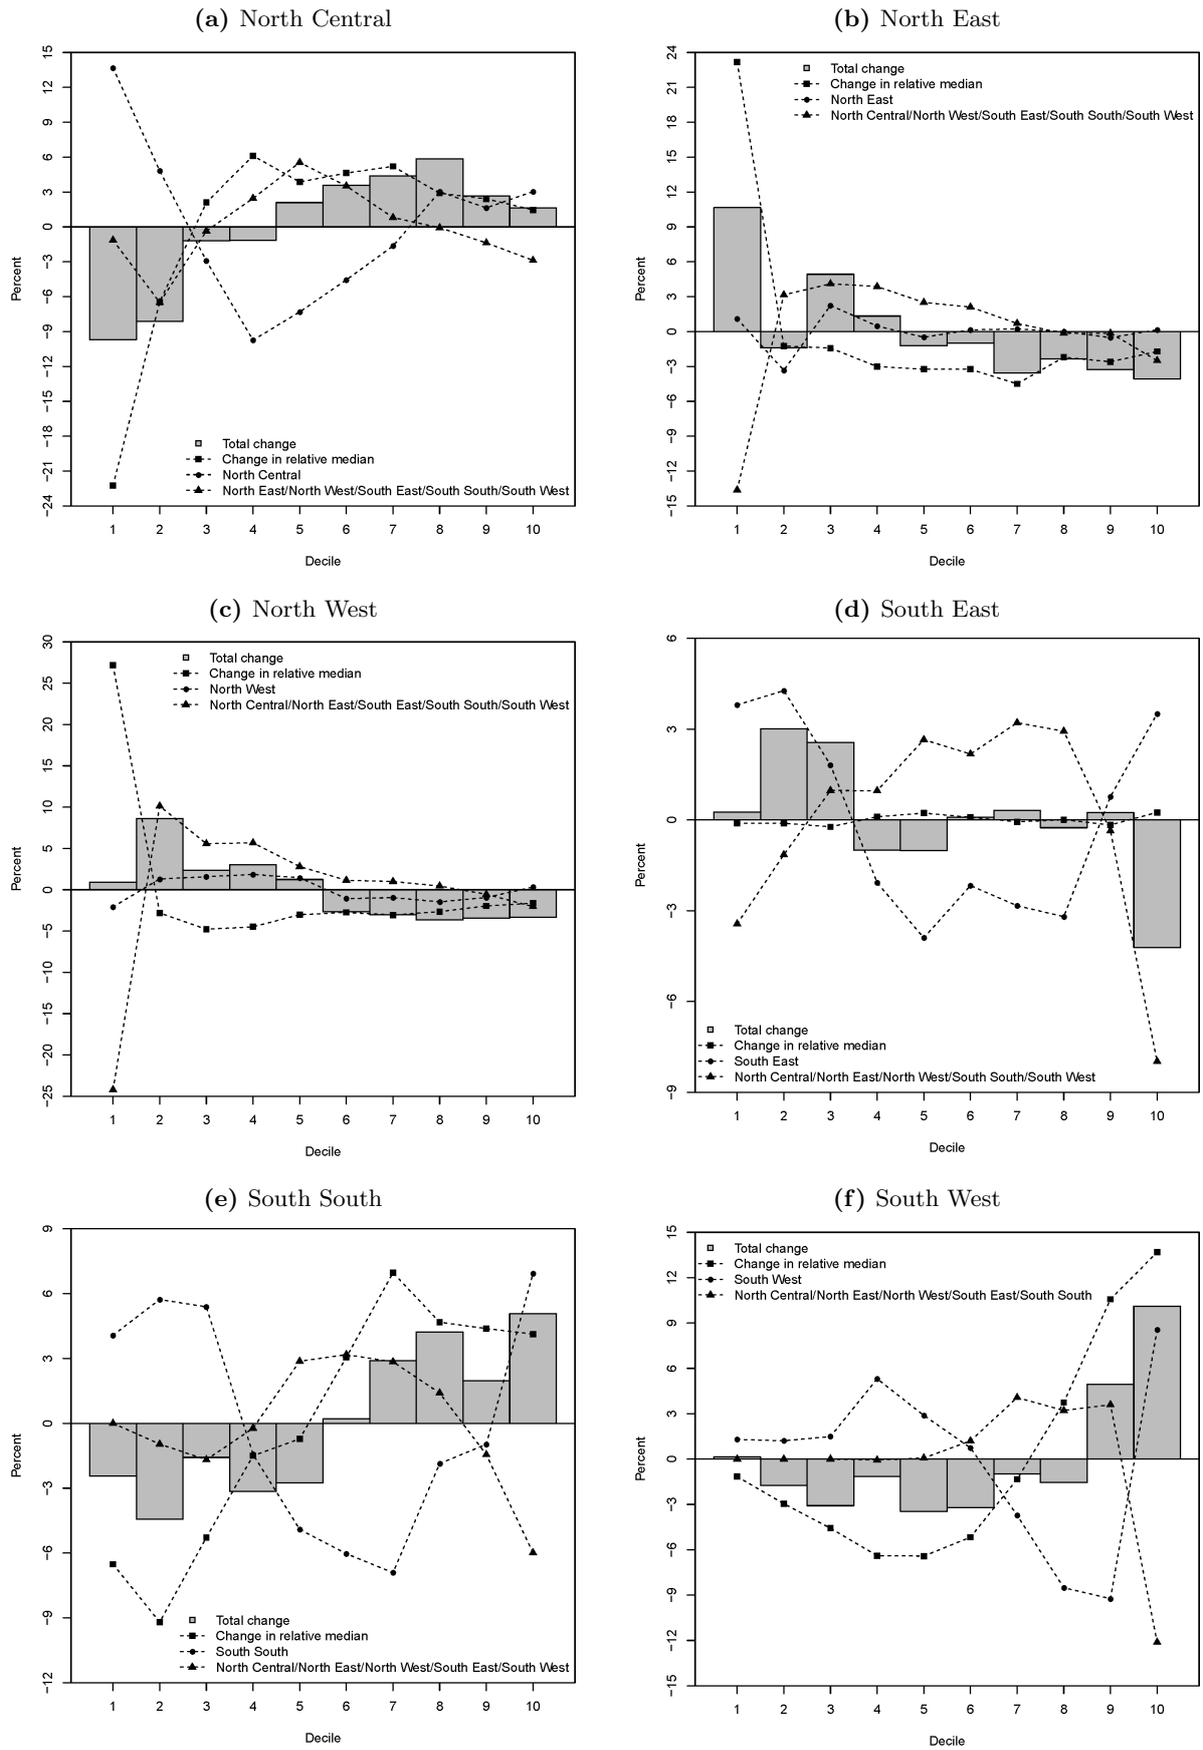


Figure 7: Sources of distributional change in the 2012/2013 to 2003/2004 relative distribution of consumption expenditures by zone.