October 2017

Measuring Poverty With Noisy and Corrected Estimates of Annual Consumption: Evidence from Nigeria

John Gibson* and Omoniyi Alimi

University of Waikato, New Zealand

Many studies of poverty use annual consumption as their welfare indicator, even though what is surveyed – particularly for food – is for a far shorter period. Extrapolating from short periods may adequately measure annual means if the sample is spread over a year but inflates measures of dispersion across households. The resulting noise in the consumption estimates will inflate measures of poverty and inequality and also leads to misclassification errors that bias logit or probit regression models of poverty determinants. In this paper we use data from the 2012/13 Nigeria General Household Survey panel to show the effect on poverty measures of using annual estimates that extrapolate from short observation periods. We derive a corrected extrapolation that does not have inflated variance. The correction method needs households to be observed in two or more times of the year, so that intra-year correlations can be estimated. Several panel surveys in Africa have recently adopted this design but no studies have yet used the corrected extrapolation method. Thus, there is an unexploited opportunity to improve the measurement and modelling of poverty in Africa.

JEL: C81, O15

Keywords: Chronic Poverty, Consumption Smoothing, Household Surveys, Transient Poverty

Acknowledgements: We are grateful to Kathleen Beegle for providing the data files. All remaining errors are those of the authors.

*Address for Correspondence: Department of Economics, Private Bag 3105, University of Waikato, Hamilton 3240 New Zealand (e-mail: jkgibson@waikato.ac.nz): 64-7-838-4289 (ph), 64-7-838-4331 (fax).

I. Introduction

Many studies of poverty use annual consumption as their welfare indicator, even though what is surveyed – particularly for food – is for a far shorter period, like a week or a fortnight. This design, with short-period observations extrapolated to annual totals, can give good estimates of annual means if the sample is spread over a year but inflates measures of dispersion across households (Deaton, 1995). The measured dispersion has both intra- and inter-household components, with some of the intra-household dispersion due to shocks that happen in the short period but are evened out over the rest of the year. Including these shocks causes the genuine between-households inequality in annual consumption to be overstated and the resulting 'noise' in the consumption estimates will inflate measures of poverty.

A failed attempt to deal with this issue was the 'usual month' design, that asked about the number of months per year that each item was acquired, the usual number of acquisitions per month, and the usual value and quantity on each acquisition occasion. This compromise between the need to measure long-run monetary welfare for each household and the need to reduce respondent burden and survey cost by using just a single interview visit was backed by Deaton and Grosh (2000) in the Living Standards Measurement Study (LSMS) manuals. Subsequent evidence shows this design cannot overcome the tension between what is feasible to ask and what it is desirable to know. The 'usual month' questions are cognitively taxing, taking far more time to ask than is needed for a fixed period recall over the same items and adding education-related inequality into measured consumption (Beegle et al, 2012), and seem to introduce errors on both the extensive and intensive margins (Friedman et al, 2017).

Instead, new guidelines for consumption surveys advocate a one-week recall for food to best balance between accuracy and cost effectiveness (FAO and World Bank, 2017). While poverty can be defined in weekly terms, to match this survey reference period, policy makers may consider such estimates less useful since the policy levers needed for pro-poor growth and pro-growth poverty reduction (Thorbecke, 2013) are usually things that act only slowly. So poverty is likely to continue to be defined over time frames like a year and extrapolation from short reference period surveys will still be used. In this paper we outline an unexploited opportunity to improve the survey measurement of poverty in several countries in Africa. In Mali, Niger, Nigeria, and Uganda, recent national panel surveys have visited the same households twice within the year, with a post-planting and a post-harvest survey round. This design is motivated by issues of agricultural measurement but it affords an opportunity – currently unused – to improve poverty measurement and modelling.

We use data from the 2012/13 Nigeria General Household Survey (GHS) panel to show the effect on poverty measures of using annual estimates that extrapolate from short observation periods. We derive a corrected extrapolation that does not have inflated variance. The average correlation between the same household's expenditures on consumption in all weeks (or months) of the year is the key parameter for correcting the estimated variance – a correlation implicitly assumed to be 1.0 by naïve extrapolation. The Nigeria GHS only lets a correlation be calculated between two periods of the year but evidence from elsewhere is that this is sufficient. For example, Gibson et al (2003) use data from urban China with two, four, or six observation periods on the same households and find the correction method works just as well with two visits, six months apart, as it does with four or six visits within the year. Notably, poverty rates estimated using the correction method matched benchmark estimates from year-long data collection, compared to a 53% overstatement of the headcount poverty rate with naïve extrapolation from a one month reference period, and a 32% upward bias if data from two, one-month reference periods, are used but the correction is not applied.

The correction method uncovers more about typical living standards because seeing the same household some months later yields new information, compared with seeing it just once or else seeing it repeatedly in short succession. Indeed, a sequence of survey visits in a short period, such as the 11 visits over a month used in Ghana, may even harm data quality due to declining compliance (Schündeln, 2017). A potentially better use of survey resources is to have fewer total visits, and to let some months elapse between each visit.

The gap between the annual poverty found with the corrected extrapolation method, and what an uncorrected short reference period survey, shows may indicate within-year transient poverty.[1] Under the "components approach" to welfare fluctuations used by Jalan and Ravallion (1998), chronic poverty is the part of total poverty due to permanent consumption being below the poverty line, while transient poverty is the remainder. The correction method gives a better measure of permanent consumption than is possible with uncorrected short-period surveys. In Nigeria it appears that half of the poverty observed in 2012/13 was within-year transient, while the remainder was chronic poverty. Likewise, the first decomposition with the corrected extrapolation method found that 50% of the poverty in Papua New Guinea was of the within-year transient type (Gibson, 2001).

Our corrected extrapolation method also has implications for the regression modelling of poverty. For per capita consumption y, poverty line z, and a vector of covariates $\boldsymbol{x}$, these models have the form: prob $[y < z \mid \boldsymbol{x}]$ with a logit or probit specification often used. This is redundant, as first noted by Ravallion (1996), since probit and logit models are based on a continuous 'latent' variable that needs to exceed some threshold for the binary outcome to occur. In contrast, consumption is fully observed rather than latent so the probability that a household is below the poverty line can be found with an OLS model for y, with weaker assumptions about the distribution of errors than probit and logit require.[2] Nevertheless, probit and logit models of poverty are popular, and six of the eight papers that use national

[1] An alternative interpretation is that the gap reflects measurement error, which is hard to empirically distinguish from transitory fluctuations without more waves of data and more econometric structure (Lee et al, 2017).

[2] For example, in *Stata* issuing the command `predict poor, pr(.,z)` after an OLS regression with y as the dependent variable would give the probability of each observation having per capita consumption below the poverty line. These calculations can be extended to poverty gap and squared poverty gap measures, and also using instrumental variables estimation, as shown by Gibson and Rozelle (2003) and Datt and Jolliffe (2005).

household survey data from Nigeria that we review in Section 2 use these models.

It is less recognized that these estimators are biased by random errors in the dependent variable, unlike for linear regression (e.g. OLS) where such errors cause imprecision but not bias (Hausman, 2001). Errors cause misclassification, where a dependent variable has a value one (e.g. a household is defined as poor) when it should have the value zero, or *vice versa*, and large biases can result. For example, Hausman et al (1988) show that with just a 5% risk of misclassification, the coefficient on a right-hand side (RHS) dummy variable (e.g. for a female-headed household) falls to 70% of its true value. At a 20% risk of misclassification the coefficient on the RHS variable is just one-third of the true value. It is even worse for a log-normally distributed RHS covariate (e.g. land holdings); with just a 2% (5%) risk of misclassification the coefficient on this variable is biased down by more than 20% (45%). Thus, probit and logit models may give misleading inferences about poverty determinants.

Since short reference period surveys provide a noisy measure of typical consumption, poverty measures based on these data will be subject to more misclassification. For example, a household whose typical consumption is above the poverty line may have a negative shock in the survey week and wrongly get classified as poor. If other households have offsetting positive shocks, a regression model of consumption will not be biased (because the error in the dependent variable is random), but a probit or logit model of poverty would yield biased and inconsistent parameters. While one solution would be to abandon logit and probit models of poverty, their enduring popularity makes this unlikely. Instead, another option is to have a less noisy measure of consumption, so as to reduce the misclassification errors in the poverty estimates, and our proposed corrected extrapolation method can assist with this.

The rest of the paper is structured as follows: Section 2 has a brief review of previous studies of poverty in Nigeria that use national samples; Section 3 derives formulae that show the degree of overstated variance when naïve extrapolation from short reference periods is used, and describes the corrected extrapolation method; Section 4 discusses the data from Nigeria used to illustrate the method, while Section 5 has the results and Section 6 concludes.

2. A Short Review of Previous Poverty Studies for Nigeria

There have been several studies of poverty in Nigeria over the last two decades, albeit with inconsistent data that hamper temporal comparisons. Some studies use the 2003/04 Nigeria Living Standards Survey (NLSS), which had seven interviewer visits to each household, at a minimum of four-day intervals in a cycle of 30 days. This was a diary-keeping survey. Such surveys may degrade into recall done when the interviewer revisits the respondent, especially if literacy rates are low (Friedman et al, 2017). The reference period was 30 days but poverty was calculated on an annual basis, with required annual food expenditure of 21,743 Naira to give 2900 calories per adult equivalent per day, plus a non-food allowance of 8,385 Naira, to give an annual poverty line of 30,128 Naira (NBS, 2005). The headcount poverty rate of 55% was an 11 percentage point decline from 1996 (a re-analysis in McKay (2013) suggests a bit less). The rural-urban gap increased; the 1996 urban poverty rate was 12 percentage points lower than the rural rate, while the urban rate was 21 points lower than the rural rate by 2004.

The next major survey was the 2010 Harmonized Nigeria Living Standards Survey (HNLSS), which showed almost no change in the poverty rate from the 2004 NLSS. There are doubts about this survey, which showed an unexplained drop in consumption in the later months that may reflect technical difficulties in survey implementation (World Bank, 2016). In light of these doubts, the World Bank used General Household Survey (GHS) panel data for 2010/11 and 2012/13 to estimate poverty. The GHS has a one-week reference period for food consumption and reference periods ranging from one-week to annual for nonfood. The baseline for temporal poverty comparisons is 2004, with NLSS data from that year having an imputed consumption figure calculated based on the 2010/11 GHS. The poverty line used was 28,830 Naira, per person per year, in 2004 prices. At this poverty line, and using the imputed consumption for 2004 and the GHS panel survey consumption for 2012/13, it seems that the national poverty rate fell from 46% in 2004 to 36% in 2013. This compares a survey that has a 30 day reference period with one that has a one-week food reference period, and both surveys extrapolate to annual totals in order to compare with annual poverty lines. Thus, the issues discussed in the current paper may matter to these poverty comparisons.

In addition to these measurement studies, several studies for Nigeria use the survey data in regression models of consumption and poverty. A selection of eight of these studies are summarized in Table 1. Most of them use NLSS or HNLSS data, for which consumption (or its sub-aggregates) is fully observed and is not a latent variable. Yet logit or probit models are used in six of these studies. Thus, the issue of estimator bias due to misclassification of poverty status, from using a short-reference period survey, is quite pertinent to the usual way that poverty is modeled in Nigeria. There are also some sub-national studies that use smaller, researcher-collected samples, which are not covered in Table 1. Some of these studies also use logit and probit models despite the criticism of this approach by Ravallion (1996).

In terms of findings from these studies, there is general agreement on the importance of location, with poverty higher for households in the North East and North West zones, all else the same. Otherwise, age, gender, education, household size, and occupation are some of the most commonly examined socio-economic correlates of poverty in these studies. There is mixed evidence for age and gender while higher education and smaller household size are generally associated with less likelihood of being poor.

3. Survey Reference Periods, Extrapolation Methods and Poverty Estimates

In order to show what assumptions are implied in forming estimates of annual consumption, and deriving poverty rates, from surveys with short reference periods, and to highlight the opportunity afforded by designs such as the Nigeria GHS that observe the same households in two periods of the year, it is helpful to consider the distribution of consumption as measured by three hypothetical surveys. One has a one-week reference period, one has a one-month reference period, and the third has an annual reference period. We abstract from issues of seasonality by assuming that the one-week and one-month reference period surveys split their samples and stagger fieldwork so that each week and month of the year are represented in the final sample with equal probability. We also abstract from any reporting errors that may affect longer reference periods due to either memory decay, if the survey uses the recall method, or to a decline in compliance, if a survey uses the diary method.[3]

The variance of the survey using the one-week reference period will exceed that of the survey with the one-month reference period. The annual reference period survey has the lowest variance. The reason is that many short term shocks tend to cancel out over a longer period. The week a household happens to be surveyed may coincide with a positive or negative income shock and, absent a full array of smoothing options (reciprocity networks, credit, and so forth), consumption in that week will be atypical. Indeed, one reason surveys historically had longer reference periods, like a fortnight or a month, was to ensure that a payday was covered in the survey period, in case spending just after a payday differed from other times. Thus, a short reference period survey, even if for a month, will have a higher variance than the true annual variance because some shocks, but not their reversal, happen in the reference period. This higher variance will overstate the share of the population who appear to live below the annual poverty line (Figure 1). The poverty gap, based on the area below the density and to the left of the poverty line, will also be inflated, as will any poverty

[3] The limited evidence on these effects is that there is a recall decay of about 3 percent for each day added to the recall period, but that it levels off at about 20 percent recall error (Scott and Amenuvegbe, 1991) while poorly supervised diaries have a similar daily decline. Eventually as the number of occasions that either consumption, purchases or acquisitions occurs exceeds some threshold that makes it too hard to recount each occasion, due either to the inherent consumption/acquisition frequency of an item or to the length of the reporting period, respondents appear to switch to rule-of-thumb estimation strategies (Gibson and Kim, 2007; Friedman et al, 2017). The reporting error then tends to level off since applying a rule of thumb to a longer or shorter period should have similar estimation errors.

measure that is distributionally-sensitive. Inequality measures are similarly inflated by the greater noise in the estimates coming from the short reference period surveys.

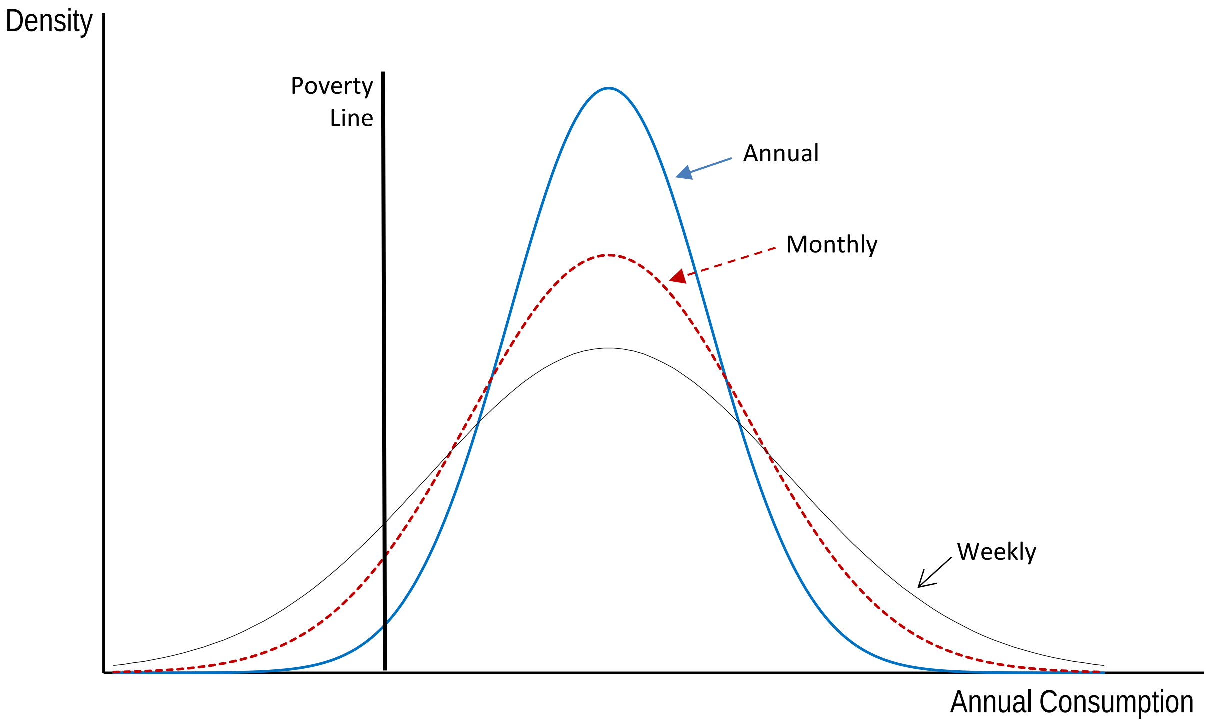

Figure 1: Excess Variability From Consumption Surveys With Short Reference Periods

It should be obvious that simply taking survey data that are based on a weekly (monthly) reference period and multiplying them by 52 (12) so as to annualize them does not alter the issue shown in Figure 1. In the rest of this section we show that the assumption that is implicit when annualizing in this way – what we call naïve extrapolation – is that the correlation between all weeks (months) in the year is 1.0 for the welfare indicator used to derive the poverty statistics. Yet the reality, shown below in the empirical part of the paper using the GHS data, is that consumption by the same household in different weeks (months) of the year can vary a lot, and this is true for both urban and rural households, and for food and non-food. Moreover, the fluctuations that drive intra-year correlation coefficients below 1.0 are seen in both levels and ranks. Thus, while these fluctuations are commonly ascribed to seasonality, they more generally reflect a reshuffling of the distribution as households are buffeted by both positive and negative shocks, from various sources at various times, that push their short-term consumption away from typical levels.

The clearest derivation is for going from a monthly period to annual consumption. We derive formulae for the excess variance, and show how intra-year correlations enable a corrected extrapolation, and then we introduce the parallel terms for weekly-to-annual. In many surveys, including the GHS, food has a weekly reference period, while non-foods have mixed reference periods ranging from one week to one year, depending on the item. Thus both correction terms that we derive, one to go from weekly to annual for food and one to go from monthly to annual for non-food, will be used in our empirical illustration.

Let $\bar{x}_m$ be the average, and $V(x_m)$ the variance, of monthly consumption expenditure across all i households and t months in the year. Extrapolating to annual totals by multiplying monthly expenditures by 12 gives an estimated variance of annual expenditures of $144 \cdot V(x_m)$. This will overstate the true variance in annual consumption expenditure, which is defined as:

(1)
$$V(x_a) = \frac{1}{N}\sum_{i=1}^{N}\left(x_{i,a} - \bar{x}_a\right)^2$$

where $x_{i,a}$ is annual expenditure for the ith household and $\bar{x}_a$ is average annual expenditures across a sample of households with N observations. Equation (1) can be expressed as:

(2)
$$V(x_a) = \sum_{t,t'=1}^{12} r_{t,t'}\sigma_t\sigma_{t'}$$

where $r_{t,t'}$ is the correlation between consumption expenditures in month t and month t' and σ_t is the standard deviation in expenditures across households in month t. This result follows because $x_{i,a} - \bar{x}_a$ in equation (1) can be expressed in terms of the sum of the deviations of each household's monthly expenditure on consumption from the mean for that month, $d_{it} = x_{it} - \bar{x}_t$ and the d_{it} terms are components of the Pearson correlation coefficient:

(3)
$$r_{t,t'} = \frac{1}{N}\sum_{i=1}^{N} d_{it}\, d_{it'} \Big/ \sigma_t\,\sigma_{t'}\,.$$

By assuming that the dispersion across households does not vary from month to month, i.e., $\sigma_t = \sigma_{t'}$, Scott (1992) shows that equation (2) can be expressed in terms of the average of all of the inter-month correlations, $\bar{r}$:

$$V(x_a) = [12 + 132 \cdot \bar{r}]\, V(x_m) . \tag{4}$$

Hence the variance from simple extrapolation to annual totals, $144 \cdot V(x_m)$, equals $V(x_a)$ only in the special case of $\bar{r} = 1$. If the average correlation is only, say, $\bar{r} = 0.6$, then the annual variance is only $91 \cdot V(x_m)$ and naïve extrapolation overstates the variance by 58%. There would also be an overstatement of any variance-based measures of poverty and inequality.

The assumption to get to equation (4), that dispersion across households does not vary from month to month, is clearly untrue. For example, heterogeneity in diet preferences has scope to have more impact in seasons of plenty (e.g., post-harvest) than in lean seasons when reduced access to food makes diets more alike (Behrman et al, 1997). Indeed, the GHS data show that the between-households variance of food consumption changes from post-planting to post-harvest. However, for this observed degree of change, assuming equi-dispersion does not seem to matter, as we show below in Figure 2. Moreover, the corrected extrapolation method yielded estimates that matched those from a benchmark of a year-long data collection in China (Gibson et al, 2003), despite the equi-dispersion property not holding there either.

The corrected extrapolation method proposed by Scott (1992) uses estimates of $\bar{r}$ to scale the *i*th household's deviation from the overall monthly average, $(x_{it} - \bar{x}_m)$ up to an annual value, and adds this to the annual average across all households, $\bar{x}_a = 12 \cdot \bar{x}_m$:

$$x_{i,A} = (x_{it} - \bar{x}_m)\sqrt{12 + 132 \cdot \bar{r}} + 12 \cdot \bar{x}_m . \tag{5}$$

For example, if the average correlation between the same household's expenditure on consumption in all pairs of months in the year is 0.5, the scaling factor is only 8.8 (=√78), rather than 12 that is implied by naïve extrapolation. Deviations of a household's one-month consumption from $\bar{x}_m$ have less effect on the annual variance than with naive extrapolation, giving a less dispersed distribution of annual consumption. Intuitively, some shocks that cause a household to deviate from the sample average for the month it was surveyed might subsequently be reversed in the rest of the year. Applying a scaling factor of 12 to these deviations, which is what naïve extrapolation does, wrongly locks them in, making it seem as if the shocks occur in each and every month of the year.

For a weekly reference period survey, the counterparts to equations (4) and (5) are:

$$V(x_a) = [52 + 2652 \cdot \bar{r}]V(x_w) \tag{6}$$

$$x_{i,A} = (x_{it} - \bar{x}_w)\sqrt{52 + 2652 \cdot \bar{r}} + 52 \cdot \bar{x}_w \tag{7}$$

where $V(x_w)$ is the variance of weekly consumption expenditures and $\bar{x}_w$ is the weekly mean. For example, if the average correlation between expenditures in all weeks of the year is 0.6, naïve extrapolation gives a variance estimate that is 1.65 times the actual annual variance, which is slightly more than the bias factor (1.58) when extrapolating from one month.

In Figure 2 we show the overstated variance for naïve extrapolation with different values for the average correlation. The bold line is for using a one-week reference period, showing by how much $2704 \cdot V(x_w)$ overstates the true variance in annual consumption, $V(x_a)$. For example, if $\bar{r} = 0.9$ the upward bias is just over ten percent, at $\bar{r} = 0.7$ the bias factor is 1.4, and it is 2.0 (3.2) if the average correlation between the same household’s consumption in all weeks in the year is as low as 0.5 (0.3). The thicker dashed line in Figure 2 is for extrapolation from a one-month reference period which gives slightly less upward bias than using a one-week reference period. The lighter dashed line relaxes the equi-dispersion assumption, letting four months of the year have a standard deviation about 30% above the other eight months (based on the post-planting and post-harvest GHS survey waves). This reduces the overstatement by about one-tenth, which is only a small adjustment, compared to the size of the overall correction factor (e.g. if $\bar{r} = 0.5$ the upward bias in the variance is still

75%), and so this complication due to $\sigma_t \neq \sigma_{t'}$ is not considered further.

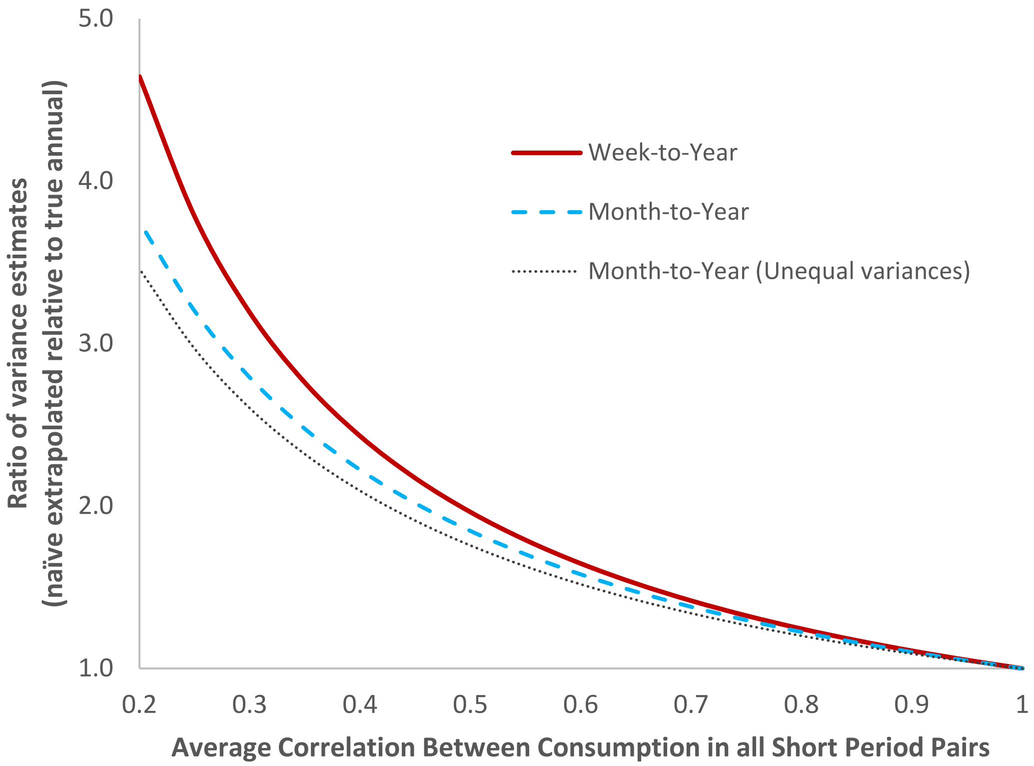

Figure 2: Variance of Short Period (or naïvely annualized) Relative to Annual Variance

The results in equations (4) and (6) clarify what is assumed by naïve extrapolation but at first glance the correction formulae in equations (5) and (7) seem to be of no practical help because compliant households would need to be surveyed every week (month) of the year, to have the data to estimate $\bar{r}$. In Ghana, where households are interviewed up to 11 times in a month, there is a monotonic decline in data quality with each successive interviewer visit, indicating reduced compliance (Schündeln, 2017). With repeated interviewing for a year the decline in compliance would likely be even worse.

Instead of needing 66 values of $r_{t,t'}$ to form the average needed for equation (5), a sampling approach can be used, estimating $\bar{r}$ from only a few of the possible inter-month correlations for the various $i \neq j$ pairs of months. This sampling approach reduces cost and burden on respondents and relies on the $r_{t,t'}$ having roughly the same value and varying little as the gap between t and t' increases. Existing evidence supports this assumption. For example, a survey from Zambia found that $r_{t,t'}$ fell by just 0.0078 for each month that the gap between t and t' increased (CSO, 1995). In the data from urban China used by Gibson et al (2003) the correlations between observations on the same household showed no statistically significant (t=0.2) time trend as the months between the two observations increased and each additional month separating observations changed the correlation coefficient by just 0.004.

The bare minimum is a survey where the same households are seen just twice within the year, approximately six months apart. This design was used by Gibson (2001) to derive corrected annual poverty estimates from a survey with a fortnight reference period for food and for most non-food. The new panel surveys in Mali, Niger, Nigeria, and Uganda have a similar design, except that the food reference period is one-week. The roughly five month gap between post-planting and post-harvest survey rounds in the GHS closely approximates two visits, six months apart, which was shown to work just as well at matching poverty estimates from a year-long survey as using four visits or six visits (Gibson et al, 2003).

4. Data and Descriptive Statistics

We use data from both rounds of the 2012/13 General Household Survey (GHS) panel, which is a subset of a larger, cross-sectional, GHS sample. A probability-proportional-to-size sample of 500 Enumeration Areas (EAs) is used, that represents the national and zonal (but not state) level. Within each selected EA, ten households are selected, to give a target sample of 5000 households. Each household is visited twice in each wave of the survey (wave one was 2010/11, which we do not use); a post-planting visit that was in September-October 2012 and a post-harvest visit that was mostly in March 2013. The survey uses a one-week recall for 105 types of food consumption (including seven types of meals out of home), and one-week, one-month, six-month and 12-month recall for 4, 29, 29, and 17 types of non-food. For the average household, non-food is only 30 percent of the total value of consumption and so we simplify the analysis by treating all of the non-food as having a one-month recall, since the major non-food items (fuels, rent, utilities, telecommunications and personal items) were all part of the expenditure groups given the one-month recall.

The consumption data that we use are from work done by the World Bank (2016) that extrapolates from the various recall periods to annual totals. After matching on the household identification variables from the post-planting and post-harvest survey rounds we have a sample of N=4565. Table 2 reports two types of correlations for consumption of the same households in these two rounds: Pearson product-moment correlations, which the correction formulae in Section 3 are based upon, and Spearman rank correlations. There are three correlations reported: for total consumption per household, for food consumption, and for expenditures on non-food consumption. The results are reported separately for the full sample, which is nationally representative, and for households in the urban and the rural sectors. If the main source of intra-year fluctuations is seasonality in food supply, we would expect rank correlations to be close to 1.0, since lowering everyone's food consumption in the hungry months of the year does not change their rank. Likewise, the non-food and urban correlations would be expected to be close to 1.0 if food production seasonality is the main reason for short-period consumption deviating from longer-term levels.

In fact, the rank correlations show there is considerable reshuffling of the distribution between the two survey rounds, which affects both food and non-food, and is apparent in the urban and rural sectors (Table 2). In the full sample, the rank correlations are only 0.79 for non-food and 0.56 for food, and they are only slightly higher for urban households. While the product-moment correlations are similar to the rank correlations for non-food, they are far lower for food, with r=0.02 in rural areas (and nationally). This difference between the two types of correlations suggests presence of some outliers in the food consumption data. These will overstate the intra-year fluctuations faced by rural households. Since our main aim is to illustrate the correction method, and noting that the data are already several years old and have been made publicly available, there is little to be gained by attempting to check and clean the data on an observation-by-observation basis. Instead, we simply trim the data by removing the top and bottom 1% of observations in terms of per capita consumption.

The removal of these potential outliers gives product-moment correlations that are closer to the rank correlations, and are more like the correlations elsewhere.[4] The intra-year correlation for food consumption in rural areas is 0.38 and in urban areas it is 0.44. The correlations for total consumption are a little higher, at 0.55 and 0.60 for rural and urban areas. Trimming the sample hardly changed the Pearson correlation for non-food (going from 0.73 to 0.71), nor did it change the rank correlations. The post-trimming stability of these correlations further suggests that the measures of food consumption may have some outliers, and attests to the need for care in handling extreme values of household survey data, since measurement error and transitory fluctuations show the same empirical patterns.

For the purposes of illustrating how naïve extrapolation overstates variances, and how the intra-year correlation enables a correction, we use correlations from the trimmed sample, of 0.38 and 0.44 for week-to-year extrapolation of surveyed food consumption in rural and urban areas, and 0.67 and 0.69 for month-to-year extrapolation for non-food consumption. With these values, the inter-household variance of annual food consumption is inflated by a factor of 2.8 in rural areas and 2.2 in urban areas, and for non-food consumption the upward bias is by a factor of 1.4, when naïvely extrapolating from short-periods to annual totals. In a summary of various intra-year correlations for household income and expenditures, and for individual labour income, McKenzie (2012) reports values that range from 0.12 to 0.79; for household expenditures observed at a time interval of six months, these averaged 0.4. Thus the correlation coefficients used here are typical of those in the (albeit limited) literature.

[4] For example, for total consumption approximately six months apart, the intra-year correlation in Papua New Guinea was 0.65, compared to 0.58 in the trimmed sample for Nigeria (and 0.06 in the untrimmed sample).

5. Results

The GHS data provided by the World Bank had annual consumption values for each survey round, from a naïve extrapolation. Averaging over the two rounds is the same as multiplying the one-week value of food consumption in each round by 26 and adding the derived figures from the two rounds (and multiplying one-month non-food consumption by six and adding). The consumption estimates from these naïve extrapolations provide the baseline estimates for comparison with the corrected values. The first three columns in Table 3 provide details on the distribution of per capita annual consumption, based on these naïve extrapolations. The overall mean is 108,300 Naira. The urban mean is 51% above the rural mean. The coefficient of variation is 0.74 in the rural sector (and nationally) and 0.65 in urban areas. The positive skew in the distribution is shown by the mean being 30% above the median.

The corrected extrapolation gives a distribution with the same means, by design, but the variance is much lower. The coefficient of variation for rural households is just 70% of what it was with the naïve extrapolation, and for urban households it is 75% of what it was.[5] The corrected extrapolation dampens the effects that any deviations from weekly or monthly means have on annual totals and this shows up in the quantiles of the annual distribution. For example, at the 10th percentile, per capita consumption is raised by about 17,000 Naira (nationally and in rural areas) by using the corrected extrapolation while at the 90th percentile it is reduced by between 23,000 Naira (rural) and 31,000 Naira (urban).

The implications of using corrected extrapolation for poverty measurement are shown in Table 4, which reports the headcount rate, poverty gap index, and poverty severity index under the two distributions.[6] We use a poverty line of 60,000 Naira, which gives a headcount

[5] In terms of Figure 2, the rural variance under naïve extrapolation is 2.04 times the corrected variance, and for urban areas the overstatement factor with naïve extrapolation is 1.76.

[6] The headcount rate is the proportion of the population in households below the poverty line, the poverty gap index is their mean distance below the poverty line as a proportion of that line, with the non-poor having a zero gap, and averaged over the full population, and the poverty severity index squares the poverty gaps in order to put more weight on the poorest.

poverty rate of 37% nationally, and 46% in rural areas and 15% in urban areas. These rates are close to those reported by the World Bank (2016), of 48% and 16% for rural and urban poverty.[7] In contrast, with the consumption estimates based on corrected extrapolation, the national headcount poverty rate is only one-half as high, at just 18%, while the rural poverty rate is 24% and the urban rate is 4.3%.

The corrected extrapolation makes even more difference to the poverty gap and poverty severity indexes, which are just one-fifth and one-tenth of their value under naïve extrapolation. The reason why these measures are more sensitive to the correction is that they depend on the mass of the density to the left of the poverty line (e.g. as in Figure 1) and that mass is greatly reduced once short-term shocks that occur during the survey period are not locked in with the implicit assumption that they occur in each and every period of the year (as the $\bar{r} = 1$ under naïve extrapolation requires). The pattern in Table 4 is consistent with the evidence from China, where poverty estimates with the corrected extrapolation matched the benchmark from year-long data collection, while poverty estimates that were based on naïve extrapolation from two, one-month reference periods six months apart, were overstated by 32% for the headcount, 78% for the poverty gap index and 112% for the poverty severity index (Gibson et al, 2003). Thus, the errors created by naïve extrapolation especially matter to measures that are sensitive to the depth and severity of poverty and they will likewise inflate measures of inequality.

The 19 percentage point gap between the estimates of the national headcount poverty rate from annual consumption estimates derived under naïve and corrected extrapolation can be interpreted as representing the within-year transitory component of poverty in Nigeria.

[7] The World Bank files have regionally deflated data but these are based on unit values (group expenditure over group quantity) from the survey, rather than from actual market price surveys. The problem with unit values is that the relative price of high-quality to low-quality items varies over space, due to transport costs, and so the within-group composition changes with these varying relative prices and the unit value cannot be a consistent indicator of the price level (Gibson and Kim, 2015). Thus, we do not use the regionally deflated series, and instead set the poverty line so as to closely match the sectoral poverty rates reported in World Bank (2016).

Under this interpretation, approximately half of the annual poverty is chronic and half is transient. The mix of these two poverty components is weighted more heavily towards the transient component in urban areas, where it is about 70% of the total cross-sectional poverty, while in rural areas the transient component is only about 47% of the total poverty. Thus, any anti-poverty interventions to offer new consumption smoothing possibilities may be expected to be beneficial to a larger proportion of the urban poor than the rural poor. More of the rural poor face chronic poverty, and for this type of poverty smoothing alone will not be enough to help them escape from their bad situation.

6. Conclusions

Survey practitioners and users of household survey data face difficult tradeoffs when trying to measure and model monetary welfare and poverty. For many purposes it is sensible to think about the long-term situation of people, and this is often measured with estimates of their annual consumption. However, it is impossible to observe a respondent for a full year, so compromise methods like the ‘usual month’ approach or having a sequence of survey visits in a short period have been used. The accumulating evidence is that these methods do more harm than good and so the thrust of recent recommendations is for consumption surveys to harmonize on a one-week food recall. However, a problem with short period surveys is that they overstate the between-households variance in annual consumption, and in poverty and inequality measures derived from these consumption data.

One potential solution that preserves the feasibility of gathering data on consumption over only a short period while still being informative about the longer-term standard of living is for a survey to make a repeated visit (or visits) to the same households some months after consumption data are first obtained. This revisit yields new information about the household, compared to seeing it just once or else seeing it repeatedly in short succession because some of the positive or negative shocks that may buffet the household when its consumption is first measured will subsequently be reversed. The key parameter for realizing the benefit of this survey design is the intra-year correlation for the multiple observations on the welfare measure. While this correlation would ideally be obtained from several within-year periods, there is some evidence to suggest that even just two visits, approximately six months apart, can enable this correlation to be reliably measured.

Although designed for other reasons, there are now several panel surveys in Africa that have the required design, with a post-planting and a post-harvest survey round that each obtain a one-week recall of food consumption and a mixed period recall for non-food for the same households. In this paper we use one of these surveys – the Nigeria 2012/13 General Household Survey – to illustrate the likely bias when short period data are annualized by naïve extrapolation (e.g. multiplying one week's food consumption by 52) and to derive a simple correction method. Based on the estimated intra-year correlations, which were only 0.38 and 0.44 for food in rural and urban areas, and 0.67 and 0.69 for non-food, we construct an adjusted annual consumption estimate that removes the excess variance. With this adjusted consumption measure, the poverty headcount rate falls to one-half of its previous value, and so this extrapolation issue is potentially important to the measurement of poverty in Africa. We caution that the data available to us may have had some errors, which acted to depress the intra-year correlations, especially for food consumption. If there is to be more use of the corrected extrapolation method that we propose, it will need to proceed hand-in-hand with greater attention to data quality and potential outliers in the consumption record since otherwise measurement error may be mistaken for transitory fluctuations in poverty.

References

Adeoti, A. 2014. Trend and determinants of multidimensional poverty in rural Nigeria. *Journal of Development and Agricultural Economics* 6(5): 220-231.

Adepoju, A, Adejare, K. 2013. Food insecurity status of rural households during the post-planting season in Nigeria. *Journal of Agriculture and Sustainability* 4(1): 16-35.

Agbaje, M, Okunmadewa, F, Oni, O, Omonona, B. 2014. Spatial Dimension of Vulnerability to Poverty in Rural Nigeria. *Quarterly Journal of International Agriculture* 53(1): 25-53.

Anyanwu, J. 2005. Rural poverty in Nigeria: Profile, determinants and exit paths. *African Development Review* 17(3): 435-460.

Anyanwu, J. 2014. Marital status, household size and poverty in Nigeria: evidence from the 2009/2010 survey data. *African Development Review* 26(1): 118-137.

Beegle K, De Weerdt J, Friedman J, Gibson J. 2012. Methods of Household Consumption Measurement through Surveys: Experimental Results from Tanzania. *Journal of Development Economics* 98(1): 19–33.

Behrman, J, Foster, A, Rosenzweig, M. 1997. The dynamics of agricultural production and the calorie-income relationship: Evidence from Pakistan. *Journal of Econometrics* 77(1): 187-207.

Central Statistical Office [CSO], 1995. *Republic of Zambia Household Budget Survey, 1993-94, Volume 1 – Main Tables and Report* Central Statistical Office, Lusaka.

Datt, G, Jolliffe, D. 2005. Poverty in Egypt: Modeling and policy simulations. *Economic Development and Cultural Change* 53(2): 327-346.

Deaton, A. 1995. 'Data and econometric tools for development analysis' in *Handbook of Development Economics, Volume III* (J. Behrman and T.N. Srinivasin, editors) Elsevier, pp. 1785-1882.

Deaton, A, Grosh, M. 2000. 'Consumption.' in *Designing Household Survey Questionnaires for Developing Countries: Lessons from 15 Years of the Living Standards Measurement Study* (M. Grosh and P. Glewwe, editors) World Bank, Washington DC, pp. 91-133.

Edoumiekumo, S, Karimo, T, Tombofa, S. 2014. Determinants of households' income poverty in the South-South geopolitical zone of Nigeria. *Journal of Studies in Social Sciences* 9(1): 101-115.

FAO And World Bank. 2017. *Food Data Collection in Household Consumption and Expenditure Surveys. Draft Guidelines for Low- and Middle-Income Countries* Discussion draft, September. Rome and Washington DC: mimeo.

Friedman, J., Beegle, K., De Weerdt, J., Gibson, J. 2017. Decomposing response errors in food consumption measurement: Implications for survey design from a randomized survey experiment in Tanzania. *Food Policy* 72(1): 94-111.

Gibson, J. 2001. Measuring chronic poverty without a panel. *Journal of Development Economics* 65(2): 243-266.

Gibson, J, Huang, J, Rozelle, S. 2003. Improving estimates of inequality and poverty from urban China's household income and expenditure survey. *Review of Income and Wealth* 49(1): 53-68.

Gibson, J, Kim, B. 2007. Measurement error in recall surveys and the relationship between household size and food demand. *American Journal of Agricultural Economics* 89(2): 473-489.

Gibson, J, Kim, B. 2015. Hicksian separability does not hold over space: Implications for the design of household surveys and price questionnaires. *Journal of Development Economics* 114(1): 34-40.

Gibson, J, Rozelle, S. 2003. Poverty and access to roads in Papua New Guinea. *Economic Development and Cultural Change* 52(1): 159-185.

Hausman, J, 2001. Mismeasured variables in econometric analysis: Problems from the right and problems from the left. *Journal of Economic Perspectives* 15(4): 57-67.

Hausman, J, Abrevaya, J, Scott-Morton, F. 1998. Misclassification of the dependent variable in a discrete-response setting. *Journal of Econometrics* 87(2): 239-269.

Jalan, J, Ravallion, M. 1998. Transient poverty in post-reform rural China. *Journal of Comparative Economics* 26(2): 338-357.

Lee, N, Ridder, G, Strauss, J. 2017. Estimation of Poverty Transition Matrices with Noisy Data. *Journal of Applied Econometrics* 32(1): 37-55.

McKay, A. 2013. Growth and poverty reduction in Africa in the last two decades: Evidence from an AERC growth-poverty project and beyond. *Journal of African Economies* 22(s1): 49-76.

McKenzie, D. 2012. Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics* 99(2): 210-221.

National Bureau of Statistics [NBS] 2005. *Poverty Profile for Nigeria*, Downloaded from: http://catalog.ihsn.org/index.php/catalog/3322/download/48492

Ogwumike, F, Ozughalu, U. 2016. Analysis of energy poverty and its implications for sustainable development in Nigeria. *Environment and Development Economics* 21(3): 273-290.

Ozughalu, U. 2016. Relationship between household food poverty and vulnerability to food poverty: Evidence from Nigeria. *Social Indicators Research* 125(2): 567-587.

Ravallion, M. 1996. Issues in measuring and modelling poverty. *The Economic Journal* 106(Sept): 1328-1343.

Schündeln, M. 2017. Multiple visits and data quality in household surveys. *Oxford Bulletin of Economics and Statistics* (forthcoming).

Scott, C. 1992. Estimation of annual expenditure from one-month cross-sectional data in a household survey. *Inter-Stat Bulletin* 8(1): 57-65.

Scott C, Amenuvegbe B. 1991. Recall loss and recall duration: An experimental study in Ghana. *Inter-Stat* 4(1): 31-55.

Thorbecke, E., 2013. The interrelationship linking growth, inequality and poverty in sub-Saharan Africa. *Journal of African Economies* 22(s1): 15-48.

World Bank. 2016. *Poverty Reduction in Nigeria in the Last Decade* World Bank, Washington DC. https://openknowledge.worldbank.org/handle/10986/25825

Table 1: Summary of Selected Previous Studies of Poverty in Nigeria that Use Data from National-Level Household Surveys

Study	Data	Estimation Method	Scope	Main findings
Ozuhalu (2016)	NLSS 2004	3-Stage Feasible Generalised Least Squares	National	Food poverty incidence was 50.2% while the vulnerability to food poverty incidence was 61.7%. Significant (but not very strong) positive correlation between food poverty incidence and vulnerability to food poverty
Ogwumike & Ozughalu (2015)	NLSS 2004	Logit	National	Male headed household, higher education of household head, residence in urban area, and higher proportion of working household members reduce the odds of being in energy poverty. Larger household size, higher age of household head, living in Northern Nigeria, and higher ratio of food expenditure to total expenditure increases the odds of being in energy poverty
Adeoti (2014)	NLSS 2004 and HNLSS 2009/2010	Logit	National; rural	Female headed households, household heads older than 60, household size (larger than 4), working in agriculture and location (North-West, North-East and South-South) increases the probability of being poor. Greater probability of being poor in 2010 relative to 2004
Anyanwu (2014)	HNLSS 2009/2010.	Logit	National	Completion of post-secondary school education, being in paid employment and residence in the North-Central and South-East zones reduces the probability of being poor. Larger household size, rural residence, no education, being a self-employed farmer, and residence in the North-West zone increases the probability of being poor.
Agbaje, Okumadewa, Oni &Omonona (2014)	NLSS 2004	3-Stage Feasible Generalized Least Squares	National; rural	Household heads with no formal education, household size, unemployment, and malaria had negative effect on consumption while household heads with tertiary education and rainfall had a positive effect on consumption in all zones. Decomposition of vulnerability shows that younger age groups are less vulnerable and larger households are likely to be poorer in the future
Edoumiekumo, Karimo & Tombofa (2014)	HNLSS 2009/2010.	Logit	South-South zone	Households headed by females, households with small family size (smaller than 5), households with larger per capita expenditure on education and health, and those with larger share of food expenditure are less likely to be poor.
Adepoju & Adejare (2013)	2010 General Household Survey-Panel. First wave Post-planting data	Probit	National; rural	Half of rural households in the country were food insecure during the post-planting period. Married household heads, larger household size and dependency ratio, living in North-Central, North-Eastern, South-East and South-West zones had significant positive effects on household food insecurity status. Female headed households, tertiary education, high expenditure on non-food items, access to both formal and informal credit and remittances had negative effects
Anyanwu (2005)	1996 National Consumer Survey dataset.	Logit	National; rural	Household size, low education, rural occupations in clerical, production and 'other' activities are positively correlated with the probability of being poor. Residence in the Central, South-East and South-South zones are negatively correlated with the probability of being poor. For male-headed households the overall patterns hold; for female-headed households, household size, education of primary school and below, and residence in the central zone of Nigeria (unlike for male-headed) are positively and significantly correlated with the probability of poverty.

Table 2: Correlation Coefficients for Between-Rounds Comparison of Total, Food, and Non-food Consumption for the Same Households, GHS 2012/13

	Full Sample (n=4565)			Trimmed Sample (n=4403)		
	Total	Food	Non-food	Total	Food	Non-food
All Nigeria						
Pearson product-moment	0.06	0.02	0.73	0.58	0.40	0.71
Spearman rank	0.68	0.56	0.79	0.69	0.57	0.79
Urban areas						
Pearson product-moment	0.18	0.09	0.79	0.60	0.44	0.69
Spearman rank	0.75	0.63	0.82	0.75	0.63	0.82
Rural areas						
Pearson product-moment	0.04	0.02	0.59	0.55	0.38	0.67
Spearman rank	0.63	0.53	0.75	0.64	0.54	0.75

Note: The trimmed sample removes the top and bottom 1% of households in terms of per capita consumption.

Table 3: Estimates of Annual per capita Consumption with Naïve and Corrected Extrapolation, GHS 2012/13

	Naïve Extrapolation			Corrected Extrapolation		
	Nigeria	Rural	Urban	Nigeria	Rural	Urban
Mean	108,300	93,730	141,560	108,300	93,730	141,560
Standard deviation	79,940	69,310	91,780	59,800	48,530	69,090
Coefficient of variation	0.738	0.739	0.648	0.552	0.518	0.488
10th percentile	40,810	37,350	57,490	57,630	54,750	78,090
Median	83,190	72,720	114,630	89,510	79,020	122,040
90th percentile	207,140	171,950	267,870	181,310	149,200	236,050

Note: Monetary values are rounded to the nearest 10 Naira.

Table 4: Estimates of Annual Poverty with Naïve and Corrected Extrapolation, GHS 2012/13

	Naïve Extrapolation			Corrected Extrapolation		
	Head-count	Poverty Gap	Poverty Severity	Head-count	Poverty Gap	Poverty Severity
All Nigeria	37.2%	9.9%	3.7%	17.5%	1.9%	0.3%
Rural areas	45.6%	12.5%	4.7%	24.0%	2.6%	0.3%
Urban areas	15.1%	3.1%	0.9%	4.3%	0.2%	0.0%

Note: The trimmed sample removes the top and bottom 1% of households in terms of per capita consumption.