

An ESRC Research Group



Consumption, welfare and well-being in Ghana in the 1990s

GPRG-WPS-002

Francis Teal

Global Poverty Research Group

Website: <http://www.gprg.org/>

The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. The work was part of the programme of the ESRC Global Poverty Research Group.

Consumption, Welfare and Well-Being in Ghana in the 1990s

GPRG-WPS-002

Francis Teal
Centre for the Study of African Economies
Department of Economics
University of Oxford

January 2005

Abstract

In this paper the links from measures of consumption to welfare and well-being are discussed drawing on a study of changes in consumption in Ghana in the 1990s. It is argued that household consumption can act as an opportunity measure of welfare. The widespread distrust of such a welfare measure may be explicable, at least in part, by how misleading it can be if averages of welfare change are presented. In Ghana it has been argued that poverty fell based on a consumption measure of welfare. In this paper it is shown that the welfare measure can be presented informatively in terms of types of households. Doing this for Ghana over the period of the 1990s reveals that while on average per capita consumption rose across all percentiles of the distribution this was not true for farmers who are on average the poorest. We also show how it is possible to present results for household types where certain characteristics remain unchanged. We focus on education and household size. The falls in welfare for farmers nearly doubles if we control for education and household size while the overall rise in consumption per capita is reduced from 11 to 3 per cent per decade. Such simple empirical analysis shows how consumption can be used as an opportunity welfare measure and that a decrease in poverty does not imply at all that most people in the economy had greater opportunities.

JEL Classification: J30 and O55.

Key words: Ghana, real incomes, poverty, well-being.

This paper draws on the Ghana Living Standard Surveys (GLSS) for the periods 1987/88, 1988/89, 1991/92 and 1998/99 which are nationally representative household surveys. The CSAE is greatly indebted to the GSO for making the GLSS data available. The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. The work was part of the programme of the ESRC Global Poverty Research Group.

1. Introduction

The utilitarian framework for the analysis of well-being has come in for some very strong criticism most notably from Amartya Sen (for example Sen (1987)). At the core of this utilitarian calculus is the notion that real income or consumption can be measured and that such money metrics are an indicator of well-being in some sense. In this paper the issues of the link from income or consumption to welfare or well-being are addressed drawing on data which provides the basis to measuring changes in household consumption over time in Ghana. Two reports have presented data showing that poverty in Ghana has declined, the first compares 1987/88 with 1991/92, the second compares 1991/92 and 1998/99, GSO (1995, 2000). It has been argued that economic reforms in sub-Saharan Africa can lead to higher incomes and lower levels of poverty, Demery and Squire (1996). The Ghana data has been prominent in this debate as Ghana is among the first of the reforming economies in sub-Saharan Africa and those reforms were sustained over the decade of the 1990s. The conclusion that poverty fell has been a source of much contention.

Why has the view that poverty fell, which sounds like a statement of fact, been so contentious? Is it that the facts are a matter of dispute or is it that the measure of poverty is thought to mislead? The poverty measure used is based on the consumption expenditure of a household. In the next section the issues involved in measuring this consumption are discussed. In section 3 we consider how far consumption can be regarded as a welfare measure drawing a distinction between outcome and opportunity measures of welfare. We ask in section 4 if consumption, viewed as an opportunity measure of welfare, has increased for most Ghanaians. We extend the comparison in section 5 by showing how growth has varied across the distribution of households. In section 6 we consider how far the rise in welfare is associated with changes in human capital and household size. A final section concludes.

2. How should consumption be measured?

The empirical question appears to be a very simple one. Have the living standards of Ghanaians improved over the decade of the 1990s? The measurement of poverty which has been followed in the case of Ghana is now one standard in the literature. GSO (2000) provides the P_1 measures for poverty due to Foster, Greer and Thorbecke (1984) broken down by location and by main economic activity. The period covered is from the third to the fourth GLSS survey (from 1991/92 to 1998/99). The study also reports a breakdown of the sources of the decline in poverty due to the “growth” effect and the “redistribution” effect using the method of Datt and Ravallion (1992), which shows that by far the most important source of the decline in poverty over this period was the growth effect. More recently the issue has arisen as to how pro-poor growth should be defined. As Kraay (2004, p. 3) notes “despite the widespread use of the term, there appears to be much less

consensus as to what exactly pro-poor growth means, let alone what its determinants are. According to one view, growth is pro-poor if the accompanying change in income distribution by itself reduces poverty (Kakwani and Pernia (2000)). However, this definition is rather restrictive, since it implies for example, China's very rapid growth and dramatic poverty reduction during the 1980s and 1990s was not pro-poor because the poor gained relatively less than the nonpoor. A broader and more intuitive definition is that growth is pro-poor if the poverty measure of interest falls. Ravallion and Chen (2003) propose this definition and apply it to a particular poverty measure, the Watts index." The welfare index which underlies the analysis of poverty in GSO (2000) is total household consumption per equivalent adult expressed in constant prices of Accra in January 1999.

This approach to poverty measurement is firmly in the "quantitative" tradition. Money metric measures of welfare add the consumption across individuals where an allowance is made for the different prices households may face. While neither measuring consumption nor allowing for differing prices are without their problems the result is a measure of real aggregate household consumption. Those sceptical of the ability of numerical analysis to capture changes in welfare have naturally been sceptical of the argument that poverty fell in Ghana. It seems rather obvious to those familiar with the poverty of farmers in Ghana and who observe the large number of young people milling around in urban centres with little, or nothing, to do that poverty has not been falling. The numbers are thought to be either wrong or misleading.

Why might the welfare measure underlying the Ghanaian quantitative analysis of the decline in poverty be misleading? The first problem is that the measure refers to households whereas we are interested in individuals. In the literature presenting money-metric measures of welfare the issue of household size has traditionally been addressed, as with the Ghana GSO (2000) data, by calculating an adult equivalent measure of household size to allow for the fact that children have lower consumption expenditures than adults, Lipton and Ravallion (1995). As noted by White (2002) the issue of the relationship of welfare to household size is an example where different disciplines appear to be able to come up with different results. "Much anthropological writing on Africa is unambiguous in the view that larger families are better off but these findings are at odds with the "stylised fact" from quantitative poverty profiles that larger households are poorer" (p.515) White discusses several possible reasons why the welfare measure may be misleading which include allowing for household economies of scale and the fact that families choose the size of the household. Thus in moving from the household, the level at which much of the data is collected, to the individual, poses considerable problems for the money-metric approach to welfare.

We propose to present the consumption measures on a per capita basis as that allows a direct comparison between the survey data for consumption and income with the macro data, differences between such sources having been an important aspect of the debate as to how fast poverty is falling (see discussion in

section 4 below). In assessing changes in welfare over time it is possible such per capita measures will be misleading if household composition has changed. We assess how household composition has changed and how those changes will affect our judgement of the money-metric as a welfare measure.

In comparing two poverty profiles over time, in the absence of panel data, it is not possible to compare the same individuals or households. Might the poverty measure be misleading as it hides important differences across different types of households? One procedure which has been developed by Ravallion and Chen (2003) is to measure the “growth incidence curve” (GIC), which shows the growth rate for each percentile of the population. Such a curve is relevant both to an assessment as to whether growth is pro-poor in the sense that the poorest percentiles of the population saw increases in their expenditures and as a valuable way of comparing welfare over time. In particular it explicitly allows for the possibility that any rise in inequality is sufficient to offset the rise in average income so the poor do not gain. Inequality matters, of course, in itself. One of the findings from the new well-being functions is that relative income is important, indeed possibly more important than the absolute level of income.

In this paper we extend this procedure by looking at the GIC for three types of households classified by the occupation of the household head: wage employees, non-agricultural (ie. urban) self-employed and farmers. In the absence of panel data this is a step closer to answering the question as to whether growth was pro-poor for all types of household. It is possible to go one step further in comparing different types of households if we are able to estimate a household level consumption function which shows how education levels and household size impact on consumption. Virtually all education is acquired before permanent entry into the labour force, so rises in household consumption due to increases in education can be thought of as an inter-generational increase in welfare. We ask how much consumption has risen for those households who saw no increases in education. Such an analysis provides insights into whether consumption is rising (and poverty is falling) because older households with given levels of education are seeing rises in welfare, an intra-generational rise, or if the welfare gain is being acquired differentially by households with higher levels of education which are likely to be the newly formed households, an inter-generational rise.

Without panel data we have no choice but to compare types of households across surveys if we wish to go beyond averages. However the use of a GIC for household types where the type is identified by the occupation of the household head provides insights into not only how poverty has been reduced but for whom.

In summary there are substantive problems that need to be tackled in showing how a measure of household consumption translates into a measure of individual welfare. Even if those problems can be solved critics of the utilitarian calculus will wish to argue that it is in no sense a proper measure of welfare. It represents a measure of the opportunities open to an individual. It explicitly ignores the use to which these

opportunities are put and their effects; as we will discuss in the next section it is to the outcomes that many argue the notion of welfare should be directed.

3. Is a measure of consumption a welfare measure?

Is the distrust of the money-metric to which reference was made above due to the conceptual problems just raised or is the objection more fundamental: that a money metric approach cannot be used to make statements about the effects on welfare and well-being? To address that issue clarity is required as to what is meant by the terms welfare and well-being.

By an increase in welfare economists working in the tradition of money-metric have meant an increase in choice, Dinwiddy and Teal (1995). Greater choice implies more opportunities which is what is meant by an increase in welfare. Such a view on welfare is directly criticised in Sen (1987, p.12): “Choice behaviour is, of course, of much interest on its own. But as an interpretation of well-being, the binary relation underlying choice is very strained. It confounds choosing with benefiting, and it does this by what looks like a definitional trick. The popularity of this view in economics may be due to a mixture of an obsessive concern with observability and a peculiar belief that choice (in particular, market choice) is the only human aspect that can be observed”. This distinction is clearly important. Being able to choose a greater range of goods is distinct from benefiting from such choice. The money metric measures of welfare are concerned with increases in choice not increases in welfare in the sense that Sen defines the term. The traditional view of welfare as increases in choice also contrasts with the notion of well-being which underlies the recent development of “happiness” functions. Such functions measure outcomes while traditional welfare measures are concerned with opportunities. The term well-being is now used [see Kingdon and Knight] to mean subjective assessments of “happiness” which is measured by means of a survey instruments which ask for an assessment of “how happy you are” on a graduated scale. It seems clear that Sen and those estimating “happiness” functions are in agreement that welfare measurement must be concerned with outcomes.

If the measure of consumption is only weakly correlated with well-being then the critics of money-measure metrics would appear to be correct in their view that it is not a good measure of welfare in the outcome sense of the term. However it seems rather doubtful that one would wish to infer, for example from a happiness function, that higher levels of consumption are not desirable because “wealth does not buy happiness”. Indeed as those working on the happiness functions recognise the most likely interpretation of the results with respect to income is that individuals adapt to some “normal” level of income. Well-being functions can be interpreted as identifying factors that determine a subjective measure of well-being in societies with given levels of income. If so they are uninformative as to how much well-being improves with

rises in consumption over time. If it turns out that a rise in consumption is only weakly correlated with welfare in the outcome sense of the term then the issue becomes how the opportunities summarised in the money-metric measure of welfare can be converted into welfare in the outcome sense of the term. A measure of consumption will be a part of any assessment of the welfare of a society if greater opportunities are regarded as desirable. The fact that it is not a sufficient statistic is clear, not simply because it measures opportunities rather than outcomes but because, as was argued above, it must be relevant who benefits from an increase in opportunity and how the spread of those opportunities changes over time.

4. The Money-Metric Measure of Consumption for Ghana

We turn now to the purely empirical, and data, issues that need to be addressed in using a consumption measure of welfare. We use Ghana as an example where, as we have already noted, consumption data has been used to argue that poverty has declined with the clear implication that welfare has risen.

The data available for assessing changes in consumption per capita in Ghana over the 1990s is presented in Table 1. We begin by setting out the figures for expenditure from the published reports on these data. Table 1 line 1 shows figures from GSO (1995) for household expenditure per capita in 1991/92 prices, line 2 shows the figures from GSO (2000) for household expenditure per adult equivalent from 1991/92 to 1998/99 in 1998/99 prices. If we link these figures to provide an index of household expenditure per capita (thus ignoring any differences between persons and adult equivalents) we obtain a rise in per capita expenditure of 35 per cent over the decade. The index number is shown in line (3) of Table 1.

In this paper we propose to adopt a measure of per capita expenditures. To that end we show, in line (4) of Table 1, the nominal figures for expenditure per capita over the four periods. It will be noted that the figures for 1991/92 and 1998/99 are very close, although not identical, to the relevant ones from the published reports. In the Table we report the CPI indices we are using and then provide two series of constant price household expenditure per capita.¹ The implications of these calculations are shown in index number form in line (9) of the Table. Per capita household expenditure rises by 16 per cent, approximately half the figure in line (3), obtained by linking the GSO studies.

Next we turn to the macro data. Table 1 lines (10)-(12) show per capita figures for GDP, investment and consumption taken from the World Bank Indicators Data for 2004. Line (13) shows the implied rise in consumption per capita to be 18 per cent which is higher than the figures from the surveys given in line (9). Line (14) shows the data from the surveys for the incomes in the principal jobs of the individuals in the labour

¹ For the third and four waves of data the CPI indices are the deflators used by the published reports. For the first two waves of the data we have linked these figures to a measure of the CPI derived from figures from the Ghana Statistical Office.

force. For reasons we will come to this income number is likely to be less reliable than the expenditure data. In the final two lines of the Table we show the logs for the per capita household expenditure and the income data. It is these which will be used as the basis of the growth rates we will present.

The first point to note from the Table is that the three sources of information available for the rise in consumption or income per capita over the period from 1987/88 to 1998/99 - macro, individual survey income based and household survey expenditure based - show rises of 18, 12 and 11 per cent respectively. There is a remarkable concordance between the individual income based data and that derived from the household expenditure data. The macro data suggest rather higher growth. It is clear that the two surveys reports should not be linked as we have done in Table 1 line (3) in part because the estimate of household expenditure per capita for round 3 of the survey was substantially reduced when the fourth round was analysed.² The published reports do not allow comparisons to be made over the decade which is our purpose in this paper.

A second rather striking finding from Table 1 is that the estimates of consumption per capita from the surveys are almost twice the numbers reported in the GDP statistics. As noted by Deaton (2003, pp3-4) assessment of how fast poverty is falling in the world depends on which statistical sources are used. "Calculations using the Penn World Tables combined with inequality measures, by Surjit Bhalla (2002) by Xavier Sala-i-Martin (2002) and by Francois Bourguignon and Christian Morrisson (2002) show rapid poverty reduction in the 1980s and 1990s....These optimistic calculations are starkly at odds with the World Bank's numbers on global poverty. The World Bank ...uses household survey data to measure directly the living standards of the poor, and their calculations show relatively little poverty reduction in the 1990s." Deaton shows that the major source for this divergence is that the PPP macro numbers show a much faster rate of growth of average consumption than do the survey based numbers. Deaton then compares survey estimates of consumption per capita with those from the national accounts for some 127 countries. He finds that "consumption estimated from the surveys is typically lower than consumption from the national accounts; the average ratio is 0.860 with a standard error of 0.029, or 0.779 (0.072) when weighted by population. (India has particularly low ratios.) The exception is sub-Saharan Africa, where the average ratio of survey to national accounts consumption is unity in the unweighted and greater than unity in the weighted calculations." (p. 7) It appears that the Ghana data is an outlier in how high are the survey estimates relative to those in the national accounts.

Figure 1 shows a comparison of the three sources of data derived in Table 1 all of which are related

² In GSO (1995, Table 2.1 p.6) the figure for consumption per capita is given as Cedis 215,000, as reported in our Table 1. At the time of the analysis of the fourth round this figure was revised down to Cedis 183,000, a reduction of 15 per cent. This figure can only be obtained from the data, not from the report, which gives figures in terms of adult equivalents rather than per capita and uses 1998 prices. In GSO (2000, p.3) there is a warning that "the results reported

to a money-metric measure of welfare change.³ The first is a measure of household per capita consumption, the second is a measure of income - both of these drawn from the GLSS survey data. The third is the data from the macro accounts. In the creation of the fourth round of the GLSS data household weights were created to address some problems in the sample design. In Figure 1 we present the data for weighted consumption. The weights were designed to be applied to households so we do not apply the weights to the individual based income data. As we have already noted it appears not to matter whether we approach the question from the basis of household per capita consumption or individual incomes.

Does this increase in average per capita consumption imply that welfare increased? We argued above that even within the context of a money-metric measure of welfare *who* benefits is a key part of its interpretation as a welfare measure in the opportunities sense of the term. In Table 2 we show how the allocation of workers across types of employment has changed across the four waves of the survey. The major change has been a shift from wage employment to urban self-employment. Thus we need to assess how far consumption and incomes of these different workers has changed.

In Figure 2 the data for income and consumption per capita are presented. In the top part of the Figure we present the rates of growth of income where individuals are classified by their principal source of income while in the bottom half we present rates of growth of household expenditure per capita where households are classified by the occupation of the household head. While both income and consumption data support the view that wage employees have done much better than farmers the perspective on the self-employed depends on whether we measure income or expenditure.

Our first problem with money metrics is apparent from this distinction between income and consumption. It is usually argued that the expenditure data is more reliable than the income, particularly for the self-employed where measurement problems are severe. If that view is accepted then it is the household expenditure data in Figure 2b which must be used as the basis for any welfare assessment. Adopting the household as the basis for measurement however creates the problem of how household averages can be linked to individuals in the household. We return to that problem below.

5 Growth Incidence Curves

In the last section we presented the average for per capita consumption by type of household. In this section we adopt the procedure used by Ravallion and Chen (2003) to present “growth incidence curves” (GIC). We begin in Figure 3a by showing the GIC for the overall average of household consumption per capita. The

here are not strictly comparable with the previous report”. In this paper we use the original data. A more detailed use of the data can be found in Teal (2001).

³ The growth rates used in the Figure are based on the differences of logarithms.

Figure shows a not dissimilar pattern to that presented by Ravallion and Chen (2003) for China in that the growth rate rises with the percentile but is positive for all. The figures differ in their magnitude. While the median rate of growth in China over the period 1990-1999 was 5.5 per cent *per annum*, the median rate of growth in Ghana was 13.3 per cent *per decade*. However growth was pro-poor in the sense that Ravallion and Chen (2003) define the term in that the growth rate was positive for the poorest percentiles.

In Figure 3b we take the next step to asking if this average figure hides important differences across types of households by presenting the GICs for the three types identified in the previous section: farmers, the urban self-employed and wage employees. It is clear that the average does hide very important differences. Growth was not pro-poor among farmers. Within the farming sector growth was negative until the 64th percentile. At the median per capita consumption fell by 3 per cent over the decade. In contrast all percentiles of self-employed and wage employee households saw rises in their per capita consumption. At the median for the self-employed and wage employees growth was 23 and 20 per cent per decade respectively.

The GICs show clearly the nature of the growth process and the process by which average consumption per capita rose. It is of course this average figure that is being cited when it is argued that poverty fell. Among the poorest percentiles the pattern of growth was very different from farmers from either the self-employed or wage employees. Poor people in these two types of household saw relatively large rises in per capita consumption while farmers saw approximately similar size falls. The net result was a modest rise in average consumption per capita as shown in Figure 3a.

6 Human Capital and Household Size

Household consumption is a function of many factors which include the composition of the household, the source of earning in the household, the number of children and its size. In the previous sections we have shown that a decomposition by household type was informative as to the patterns by which consumption changed over the decade of the 1990s. In this section we ask if it is informative to focus on two aspects of the household which are known to be important - its level of human capital and its size - and how doing so affects our interpretation of the money-metric measure of welfare. We do this by controlling for the education of the household head and the size of the household measured by the log of the total number of household members. With such controls the increases in per capita consumption have an interpretation as to how much consumption would have risen for a household with given levels of education and of given size. We now add these two sets of controls - for education and household size - to each of the three types of household.

In Figure 4a we control simply for human capital, in Figure 4b we control for both human capital and household size. The rise in decadal consumption per capita shown in Figure 1 is 11 per cent. This is reduced

to 8 per cent with controls for human capital and to 3 per cent with controls for both human capital and household size. The relatively modest reduction due to controls for human capital suggests that the welfare gain is mainly an intra-generational rather than inter-generational gain. The effect of the controls for household size is, across all three types of household, to reduce the underlying rise in consumption although these effects are much greater for farmers and the self-employed than for wage employees. While these results are in no sense causal, they do not tell us what happens to household consumption if we alter education or size, they do however tell us what on average happened to types of households over this ten year period. For households of a given size, and given level of education, household consumption has risen by only 3 per cent over the decade. Further, and more important from a welfare perspective, they show that the differences across household type are further accentuated by these controls. Real expenditure per capita for farmers fell by 7 per cent. We know from the GIC of the previous section that most farmers saw falls in their expenditures. For those households who saw no change in size or increase in human capital this fall will have been much greater.

7. Conclusions

Have the living standards of Ghanaians improved over the decade of the 1990s? The answer is clear: it all depends who is being referred to. The largest and poorest section of Ghanaians, the farmers, saw a fall in their expenditures over the decade of some 3 per cent. In contrast the urban self-employed and wage employees saw rises in excess of 17 per cent. Households headed by wage employees or those with urban self-employment have close to twice the per capita expenditure of farmers. The average consumption measure, which rose by 11 per cent, is what underlies any assertion that poverty fell. The reason the assertion has proved so contentious is clear: such averages hide falls in consumption amongst the poorest and rises for others.

One of the advantages of the money-metric measure of welfare is that it is possible to ask what happened to different types of household. We have carried out two types of analysis. The first has looked at GIC such as those used by Ravallion and Chen (2003) to assess if the poor benefited from growth. The second has looked at what happened to per capita consumption for those with given levels of human capital and without changes in household size.

Calculating a GIC for Ghana over the period of the 1990s reveals that while on average per capita consumption rose across all percentiles of the distribution this was not true for farmers who are on average the poorest. While on average growth was pro-poor within the poor there were substantial gainers and losers. The controls for human capital and household size do not alter the pattern by which farmers saw falls in their

average per capita consumption while other groups saw rises but they very substantially change the magnitude. On average there was a rise of only 3 per cent in per capita consumption for all households when we control for human capital and household size. For these types of households we see a pattern by which the best off households, those headed by a wage employee, saw the largest rise while farming households saw falls of 7 per cent. While such controls have no causal implications they do provide valuable descriptive statistics.

The data from Ghana has been used to show how a money-metric measure of consumption can be linked to well-being defined in terms of the opportunities made available to the household. Such an approach is quite different from one based on some measure of happiness (as in “happiness” functions) or any outcome based measure of welfare (such as Sen argues should be used). Such a welfare measure is explicitly concerned not with outcomes but with opportunities. In so far as greater opportunities are thought desirable then we would argue that the money-metric can be used as a welfare measure. The empirical analysis has shown that it can be used to show that a decrease in poverty does not imply at all that most people in the economy had greater opportunities.

