The African Growth Miracle

Abstract

Measures of real consumption based upon the ownership of durable goods, the quality of housing, the health and mortality of children, the education of youth and the allocation of female time in the household indicate that sub-Saharan living standards have, for the past two decades, been growing about 3.5 percent per annum, i.e. three and half to four times the rate indicated in international data sets.

Alwyn Young
Department of Economics
London School of Economics
This Draft: February 2010

I am grateful to Chad Jones, Pete Klenow and seminar participants at the NBER Productivity Group, NYU and Yale for helpful comments.

I. Introduction

Much of our current understanding of the factors behind growth and development, and our continuing attempts to deepen that understanding, is based upon cross-national estimates of levels and growth rates of real standards of living. Unfortunately, for many of the poorest regions of the world the underlying data supporting existing estimates of living standards is minimal or, in fact, nonexistent. Thus, for example, while the popular Penn World Tables purchasing power parity data set version 6.1 provides real income estimates for 45 sub-Saharan African countries, in 24 of those countries there has actually never been any benchmark study of prices.[1] In a similar vein, although the on-line United Nations National Accounts database provides GDP data in current and constant prices for 47 sub-Saharan countries for each year from 1991 to 2004, the UN statistical office which publishes these figures had, as of mid-2006, actually only received data for just under half of these 1410 observations and had, in fact, received no constant price data, whatsoever, on any year for 15 of the countries for which the complete 1991-2004 on-line time series are published.[2]

Where official national data are available for developing countries, fundamental problems of measurement produce a considerable amount of unquantified uncertainty. As noted by Heston (1994), consumption measures for most developing countries are derived as a

[1]See "Data Appendix for a Space-Time System of National Accounts: Penn World Table 6.1", February 2008. As explained in the source, expatriate post-allowance indices are used to extrapolate the price studies of benchmark countries to non-benchmark economies.

[2]This statement is based upon a purchase in 2006 of all the national accounts data records ever provided to the UN Statistics Division by member countries. When queried about the discrepancy between the completeness of their website and the data I had purchased, UN officials were quite frank about the difficulties imposed by the demands from users for a complete series, and their website openly explains that much of their data is drawn from other international organizations and extrapolations (http://unstats.un.org/unsd/snaama/metasearch.asp). Similar frankness concerning the need to use extrapolations from the data of other countries to fill in gaps is present on the World Bank data website (see http://go.worldbank.org/FZ43ELUKR0).

residual, after subtracting the other major components of expenditure from production side estimates of GDP. Production side estimates of subsistence and informal production and other untaxed activities are, however, very poor, leading to gross errors in the calculation of consumption levels. Thus, for example, the first national survey of the informal sector in Mozambique in 2004 led to a doubling of the GDP estimate of nominal private consumption expenditure. Where direct surveys of consumer expenditure are available in developing countries, these must also be treated with care, given the difficulty of collecting accurate nominal consumption data. This is best illustrated by the case of the United States where the considerable difference between the growth of reported expenditure in the Consumer Expenditure Survey and the NIPA (using the production residual method) led to about a ln 40 percent gap between the two series by the early 1990s (Slesnick 1998). The problems of getting accurate reports of household expenditure, and marrying them to appropriate price indices, should be even greater in poor countries with limited resources devoted to collecting data from individuals with minimal education.

The paucity and poor quality of living standard data for less developed countries is well known and is motivating expanding efforts to improve the quality of information, as represented by the World Bank's International Comparison Programme and Living Standards Measurement studies. However, there already exists, at the present time, a large body of unexamined current and historical data on living standards in developing countries, collected as part of the Demographic and Health Survey (DHS). For more than two decades this survey has collected information on the ownership of durables, the quality of housing, the health and mortality of children, the education of the youth and the allocation of women's time in the home and the market in the poorest regions of the world.

In this paper I use the DHS data to construct estimates of the level and growth of real consumption in 29 sub-Saharan and 27 other developing countries. These estimates have the virtue of being based upon a methodologically consistent source of information for a large sample of poor economies. Rather than attempting to measure total nominal consumption and marry it to independently collected price indices, they employ direct physical measures of real consumption that, by their simplicity and patent obviousness (the ownership of a car or bicycle, the material of a floor, the birth, death or illness of a child), minimize the technical demands of the survey. While the items they cover provide little information on comparative living standards in developed countries, in the poorest regions of the world they are clear indicators of material well being, varying dramatically by socioeconomic status and covering, through durables, health & nutrition and family time, the majority of household expenditure.

Econometrically, my procedure recognizes the error involved in product sampling by, first, placing greater weight on products which have a stronger statistical correlation with real incomes and, second, discounting the observations in product groups which are strongly autocorrelated. Where estimates are constructed in multiple steps, I explicitly take into account the estimation error of earlier steps and its effect on the inferences that can be drawn in later calculations. As I make direct use of micro data, my estimates incorporate an evaluation of the uncertainty introduced by sample sizes and the reduced information implied by correlation within clusters. In sum, I estimate the level and growth of living standards in 56 developing economies using a methodologically consistent information set of easily measurable and economically significant indices of material well being employing econometric techniques that take account of correlation within clusters and product groups and biases and uncertainties produced by multi-step procedures to produce point estimates and their associated standard errors.

The principal result of this paper is that real household consumption in sub-Saharan Africa is growing between 3.2 and 3.8 percent per annum, i.e. three and a half to four times the 0.9 to 1.0 percent reported in international data sources. This growth is not due to the influence of any particular product group, as durables, housing, health, and family economics all show growth which is at least double that reported in international sources. The growth of non-African economies is also higher than reported in international sources, but the discrepancy here is much less pronounced, with growth of 3.1 to 3.8 percent, as opposed to the 1.7 to 2.2 percent indicated by international sources. While international data sources indicate that sub-Saharan Africa is progressing at less than half the rate of other developing countries, the DHS suggest that African growth is easily on par with that being experienced by other economies. Regarding the cross-national dispersion of real consumption, the DHS data suggest levels that are broadly consistent and highly correlated with those indicated by Penn World Tables, although there are substantial differences for individual countries.

I present my methodology and results in stages, allowing the reader to more easily absorb the different components that make up the approach and also establishing, I hope, that the basic result concerning sub-Saharan growth is extremely robust. I begin, in section II, by describing the DHS data and the durable goods, housing, nutrition & health, and household time measures of real consumption. Section III then presents an introduction to my methodology, showing how aggregate product level consumption data provides information on the ratio of the growth rate to the cross national standard deviation of real living standards. Intuitively, the trend and cross national dispersion of product level consumption is related, through the income elasticity of demand, to the trend and dispersion of aggregate real consumption expenditure so that the ratio of these two measures for a sample of products provides information, once one adjusts for product specific idiosyncratic effects, on the equivalent ratio for real living standards. Section IV implements this idea, highlighting the gross inconsistency between the DHS and the most popular of real living standards measures, the Penn World Tables. Put simply, the DHS data imply a ratio of growth to cross sectional dispersion on the order of 2 (non-Africa) to 4 (Africa) times that present in the PWT. Either the cross sectional dispersion or the growth rates of real consumption expenditure implied by the two sources are radically different.

Section V continues by showing how the use of the micro correlation between educational attainment and consumption levels present in household datasets allows one to infer the income elasticity of demand for each of my real consumption measures, thereby allowing the separate estimation of both the growth rate and standard deviation of real living standards. When implemented using the DHS data in section VI, these methods indicate that the discrepancy between the DHS and the PWT lies in the growth rate of real consumption (as noted above) and not its cross national dispersion, which is roughly equivalent in the two sources. To provide a robustness check on the results, section VII extends the methodology to allow for heterogeneous (local) demand patterns. In conventional national income accounting the use of common global prices allows a comparison of both the levels and growth rates of living standards, but the use of local prices restricts international comparisons to growth rates. Similarly, the movement from the assumption of a homogenous (global) pattern of demand in sections V and VI to allowing local demand patterns in section VII restricts the analysis to growth alone. Nevertheless, this extension provides an importance robustness check as it allows for the possibility that, for the products I examine, African income elasticities of demand are uniformly higher, so that the observed movements in product level consumption are associated with smaller overall consumption expenditure growth. When implemented, in section VIII, this methodology produces, if anything, even higher estimates of African consumption growth.

As a final attempt at reconciliation, section IX notes that conventional aggregate data such as the PWT report the ln of mean consumption, while my micro-data based DHS calculations actually concern the mean of ln consumption, the difference between the two measures representing the degree of consumption inequality. Simple recalculations allow me to convert my measures to ln of the mean equivalents. These lower the estimated growth of consumption somewhat, as faster rural growth is lowering inequality in both the African and non-African economies. With locally estimated income elasticities, I find that the ln of mean consumption is growing 3.5 percent in sub-Saharan Africa and 3.1 percent in the non-African economies. Section X concludes with figures which show the broad agreement between the DHS and PWT regarding levels of consumption and the immense gap in their assessment of the absolute and relative (to other countries) growth of sub-Saharan African living standards. An appendix details the construction and coding of variables drawn from the DHS surveys.

II. Demographic and Health Survey Data on Living Standards

The Demographic Health Survey and its predecessor the World Fertility Survey, both supported by the U.S. Agency for International Development, have conducted irregular, but in-depth, household level surveys of fertility and health in developing countries since the late-1970s. Over time the questions and topics in the surveys have evolved and their coverage has changed, with household and adult male question modules added to a central female module, whose coverage, in turn, has expanded from ever married women to all adult women. I take

1990 as my starting point, as from that point on virtually all surveys include a fairly consistent household module with data on household educational characteristics and material living conditions that are central to my approach. In all, I have access to 135 surveys covering 56 developing countries, as listed in Appendix I. The sample consists of about 0.5 million households in sub-Saharan Africa and 1.1 million households in the rest of the world, including useful information on 3.1 million youths aged 6-24, 1.2 million currently married women aged 15-49, and 0.65 million children less than 3 years of age.[3]

The raw data files of the DHS surveys are distributed as standardized "recode" files. Unfortunately, this standardization and recoding has been performed, over the years, by different individuals using diverse methodologies and making their own, idiosyncratic, errors. This produces senseless variation across surveys as, to cite two examples, individuals with the same educational attainment are coded as having dramatically different years of education or individuals who were not asked education attendance questions are coded, in some surveys only, as not attending. In addition, there are underlying differences in the coverage of the surveys (e.g. children less than 5 years vs. children less than 3 years) and the phrasing and number of questions on particular topics (e.g. employment) which produce further variation. Working with the original questionnaires and supplementary uncoded raw data generously provided by DHS programmers, I have recoded all of the individual educational attainment data, corrected coding errors in some individual items, recoded variables to standardized definitions and, as necessary, restricted the coverage to a consistent sample (e.g. married women, children less than 3 years)

[3]These numbers represent the sample with relevant information. Thus, for example, the 1.6 million households actually contain 3.5 million youths, 2 million women, and .9 million children. Youth school attendance is not collected in some surveys and, depending upon the survey, women are interviewed in depth (providing information on themselves and their young children) according to whether they have ever been married and/or slept in the house the previous night or are usual members.

and removed surveys with inconsistent question formats (in particular, regarding labour force participation). Appendix I lists the details.

I use the DHS data to derive 26 measures of real consumption distributed across four areas: (1) ownership of durables; (2) housing conditions; (3) children's nutrition and health; and (4) household time and family economics. Table I below details the individual variables and sample means. All of these variables are related to household demand and expenditure, broadly construed, and, as shown later, are significantly correlated with real household incomes, as measured by average adult educational attainment. I have selected these variables on the basis of their availability and with an eye to providing a sampling of consumption expenditures that would, through material durables, nutrition & health and household time, cover most of the budget of households in the developing world. I have made the decision to break measures of household time into different age groups to account for different demand patterns at different ages as the possibilities of substitution between home production, human capital accumulation and market labour evolve. Thus, for example, in richer households young women are more likely to be in school and less likely to be working in the late schooling years (ages 15-24), but, consequently, are more likely to be working as young adults (ages 25-49). Although males are included in the schooling and children's health variables, I do not include separate time allocation measures for adult males because male questionnaire modules are less consistently available and male participation behavior, when recorded, is less strongly related to household income and, hence, by my methodology, would play little role in estimating relative living standards.

My approach will be to use the correlation between real consumption in a sample of products and relative household incomes, as measured by adult educational attainment, to draw inferences about levels and trends in relative regional consumption expenditure. While the full

Table I: DHS Real Living Standard Measures by Category					
	N	Mean		N	Mean
Ownership of Durables			Housing Conditions		
Radio	1557550	.574	Electricity	1534362	.528
Television	1577616	.405	Tap Drinking Water	1569114	.451
Refrigerator	1473490	.249	Flush Toilet	1449330	.322
Bicycle	1489805	.296	Constructed Floor	1400359	.598
Motorcycle	1431210	.102	ln # Sleeping Rooms per Person	717178	-.927
Car	1460012	.066			
Telephone	1130847	.172			
Children's Nutrition and Health			Household Time and Family Economics		
ln Weight (100g)	465085	4.44	Attending School (age 6-14)	1916473	.712
ln Height (mm)	454582	6.59	Attending School (age 15-24)	1219551	.304
No Diarrhea	590540	.799	Working (women age 15-24)	195060	.416
No Fever	578304	.676	Working (women age 25-49)	588049	.554
No Cough	582544	.658	Gave Birth Past Year (age 15-24)	289763	.312
Alive	649386	.930	Gave Birth Past Year (age 25-49)	898526	.141
			Ever Married (women age 15-24)	726630	.431
			Ever Married (women age 25-49)	1083877	.936

Notes: All variables, other than ln weight, height and rooms per capita, coded as 0/1. Ownership of Durables: at least one such item in the household; Housing Conditions: constructed floor means made of other than dirt, sand or dung. Household Time: individual variables, i.e. coded separately for each individual of that age in the household; recent fertility and market participation refer to currently married women only. Children's Health: individually coded for each child born within 35 months of the survey; diarrhea, cough and fever refering to the absence of these for the individual in question (if alive) in the preceding two weeks; ln weight and ln height referring to measurements of living children at the time of the survey.

methodology is discussed in later sections, a few obvious concerns should be noted at this time. First, there is likely to be a significant covariance, independent of income, between many of these variables at the national or regional level, so that they cannot be construed as a true, independently drawn, random sample of consumption levels. The use of random effects within the four broad product groups (e.g. housing), in both trends and within regions, will address this issue, discounting the number of observations to the degree that there are strong empirical correlations within but not across groups, i.e. to the degree that the product groups represent correlated demand along narrow dimensions. Second, there are likely to be regional, idiosyncratic, factors affecting the measured levels of many of these variables, independent of relative consumption expenditure. Again, these can be explicitly recognized with random effects at the product and product group level, identifying the degree to which there is correlation in levels within products and product groups that do not extend across all products, producing more efficient estimates of overall country real consumption levels with appropriate standard errors. Third, local, district level, infrastructure is an important determinant of the realized household consumption of some products. Arguably, individuals choose their residence precisely to get access to such communal infrastructure, and pay for it implicitly through the local cost of housing and land, so it should be viewed as an element of demand. Nevertheless, I will address this issue by estimating demand equations using cluster level random and fixed effects. The results are not dramatically different. Finally, a skeptic might question whether the consumption of some of these "products" is even related to real expenditure. As presented later, my methodology is inherently "idiot proof", as products whose micro level correlation with real incomes is poor will play an insignificant role in determining the estimated growth and levels of regional consumption.

As a final comment on the data, I should note that the DHS codes households as living in urban (cities and towns) or rural (countryside) areas. In what follows, I estimate average levels of demand at the urban and rural level. To calculate and report national averages, I divide the total DHS urban and rural household weights by their combined sum to arrive at urban/rural household shares, which I then use to weight the estimated urban/rural consumption levels to produce national totals. For the 56 countries across the 135 surveys in my sample, the urban share varies from a minimum of .06 in Rwanda 1992 to .80 in Brazil 1996, with a mean of .38 and standard deviation of .18. There is some apparent random variation in the DHS's estimate of urban/rural shares, but the trend in average country-level urbanization, at .4 of one percent per annum within and outside Africa, does not seem unreasonable. With an urban-rural ln consumption gap of .80 outside of Africa and 1.03 in sub-Saharan Africa, this trend contributes no more than .3 to .4 percent of growth. I also report, in section IX, separate, unweighted, estimates for urban and rural regions.

III. Methods I: Using Consumption Aggregates

I introduce my methodology in stages: first, in this section, showing how the ratio of the growth to the standard deviation of living standards can be inferred from panel data on product level consumption aggregates, then, in a later section, showing how the within survey micro-level correlation between product consumption and educational attainment can be exploited to produce separate point estimates of both the growth and standard deviation of living standards and, yet further on, showing how international heterogeneity in demand patterns can be accommodated while still producing internationally comparable measures of real consumption growth. Proceeding in this sequential fashion makes the pieces that make up the overall methodology transparent and digestible and, as the reader will see, highlights the inconsistency between the DHS data and the most recognized of international measures of living standards, the Penn World Tables.

Let some measure of the demand for product p be given by:

$$(1)\quad \ln(Q_p) = \alpha_p + \eta_p \ln(C^N) + \vec{\xi}_p' \ln(\vec{P})$$

where α_p is a constant, η_p the quasi-income elasticity of demand, C^N nominal consumption expenditure, $\vec{\xi}_p'$ a vector of own and cross quasi-price elasticities of demand, and $\ln(\vec{P})$ the associated vector of prices relative to some base. I use the term quasi in describing the elasticities, because ln(Q) need not be actual ln quantity demanded, but only some measure related to that quantity, such as the index in a probability model or an outcome of food demand such as body weight. Homogeneity of demand of degree 0 in expenditure and prices implies that that the quasi-income elasticity of demand equals the negative of the sum of the own & cross quasi-price elasticities:

$$(2)\ \ \eta_p = -\sum_q \xi_{pq}$$

Equation (2) holds even when Q is not strictly speaking quantity demanded, as anything associated with that demand should, equally, have the same homogeneity of degree 0 property.

To reformulate (1) in terms of real consumption, we add and subtract from nominal expenditure the expenditure share weighted movement of prices from the base to produce

$$(3)\ \ \ln(Q_p) = \alpha_p + \eta_p[\ln(C^N) - \vec{\Theta}'\ln(\vec{P})] + \eta_p[\vec{\Theta}' + \vec{\xi}_p'/\eta_p]\ln(\vec{P})$$

where $\vec{\Theta}$ is a vector of product expenditure shares.[4] I operationalize (3) empirically by taking the last term on the right hand side as the error term:

$$(4)\ \ \ln(Q_p) = \alpha_p + \eta_p \ln(C^R) + \eta_p \varepsilon^{\vec{P}}$$

where the superscript $\vec{P}$ on ε is used to emphasize the role relative prices play in determining the error term. Clearly, $\vec{\Theta}$ and $\vec{\xi}_p/\eta_p$ are vectors whose components sum to one and negative one, respectively, so that when added they sum to zero. Consequently, uniform inflation drops out of the error term which, when normalized by the quasi-income elasticity, is a zero-weight average of relative price changes; something that, arguably, is homoskedastic across products and has an expected value of zero.

[4]These are actual product expenditure shares, and not quasi in any way, but, as will be seen, there is no need to actually ever compute them.

Real consumption expenditure is growing, worldwide, at an average rate g, so that real consumption per capita in country c at time t can be written as:

$$(5)\quad \ln(C^R_{ct}) = \ln(C^R_c) + g*t + g_c*t$$

where g_c represents the deviation of the country's growth rate from the global average g and $\ln(C^R_c)$ equals ln relative consumption in the base year. Substituting (5) in (4) we see that the data generating process is given by:

$$(6)\quad \ln(Q_{pct}) = \alpha_p + \eta_p \ln(C^R_c) + \eta_p g*t + \eta_p g_c*t + \eta_p \varepsilon^{\bar{P}}_{pct}$$

Consider running for a single product p a random effects panel regression on a sample of consumption data in countries c at times t:

$$(7)\quad \ln(Q_{pct}) = c_{pc} + g_p*t + v_{pc}*t + e_{pct}$$

where, with the subscripts p used to remind us that each regression is specific to a particular product, c_{pc} represents a set of country level dummies, g_p the average product consumption growth rate, v_{pc} a random effect accounting for variation in country growth rates, and e_{pct} the putative iid error term. Comparing (7) and (6), one sees that if e_{pct} is truly iid then:

$$(8)\quad \hat{c}_{pc} = \alpha_p + \eta_p \ln(C^R_c) \qquad \hat{g}_p = \eta_p$$

so that, with σ[x] denoting the standard deviation of x:

$$(9)\quad \frac{\hat{g}_p}{\sigma[\hat{c}_{pc}]} = \frac{\eta_p g}{\eta_p \sigma[\ln(C^R_c)]} = \frac{g}{\sigma[\ln(C^R_c)]}$$

i.e., the ratio of the time-series to cross-sectional variation of product consumption levels gives information on the ratio of the growth rate to the underlying standard deviation of real consumption.[5]

[5]Two points deserve mentioning. First, if the unobserved income elasticity of demand is negative, what (8) and (9) report as positive growth actually represents declining real expenditure. This is not a

Unfortunately, e_{pct} is not likely to be iid, as there are persistent influences on demand within a country other than real expenditure, most notably sustained relative price differences. Hence, the underlying data are actually better described as being produced by the process

$$(6)' \quad \ln(Q_{pct}) = \alpha_p + \eta_p \ln(C_c^R) + \eta_p g * t + \eta_p g_c * t + \eta_p \varepsilon_{pc}^{\bar{P}} + \eta_p \varepsilon_{pct}^{\bar{P}}$$

where $\varepsilon_{pc}^{\bar{P}}$ represents a persistent country x product error and $\varepsilon_{pct}^{\bar{P}}$ an iid residual error. In this case, the product x country constants include the influence of relative prices on product demand so that

$$(8)' \quad \hat{c}_{pc} = \alpha_p + \eta_p \ln(C_c^R) + \eta_p \varepsilon_{pc}^{\bar{P}} \qquad \hat{g}_p = \eta_p$$

and (9) actually estimates

$$(9)' \quad \frac{\hat{g}_p}{\sigma[\hat{c}_{pc}]} = \frac{g}{\sigma[\ln(C_c^R) + \varepsilon_{pc}^{\bar{P}}]}$$

It seems likely that there is a fair amount of independent variation in $\varepsilon_{pc}^{\bar{P}}$ so that, on average, (9)´ will understate the ratio of growth to the standard deviation of real living standards.

The problem of persistent, non-expenditure related, country specific influences on the demand for individual products can be solved by extending the single product random effects regression to a sample of products:

$$(10) \quad \ln(Q_{pct}) = a_p + b_p(c_c + g * t + v_c * t + u_{pc} + e_{pct})$$

---

problem in later sections, where I estimate the income elasticity of demand, and I finesse it in this section by measuring each variable so that (as confirmed later) an increase is positively associated with higher real incomes, e.g. women *not* giving birth, *not* being married, etc.

Second, as all of this, and everything to come, is maximum likelihood, all statements about coefficients being equal to parameters are actually statements about asymptotic consistency. I generally have about 130+ observations at the product level regression presented so far, and close to 3200 when I combine all products to produce my final estimates (discussed further below). For the micro-level analysis presented in later sections, there are typically at least 1/2 million observations, with thousands of observations per regional dummy.

where all error terms are uncorrelated with each other and across subscript categories, e.g. $E(u_{ic},u_{jc}) = 0$ $(i \neq j)$. With data from a number of products, product level persistent differences in consumption u_{pc} can now be separated from overall relative levels of consumption c_c which, along with the average growth rate g, are estimated by the cross-sectional and time-series co-movement across products of consumption, intermediated, in magnitude, by the quasi-income elasticities b_p. Given the multiplicative way they enter with the other coefficients in the regression, the quasi-income elasticities are only identified in a relative sense, and identification is achieved by restricting one of them, say that for the first product, to equal 1.[6] In this case

$$(11)\quad \hat{c}_c = \eta_p \ln(C_c^R)/\eta_1 \qquad \hat{g} = \eta_p/\eta_1$$

Although the actual growth and cross sectional dispersion of real living standards are still not known, the ratio of the two is purged of product level variation in consumption levels:

$$(12)\quad \frac{\hat{g}}{\sigma[\hat{c}_c]} = \frac{g}{\sigma[\ln(C_c^R)]}$$

It is obvious that there are various generalized least squares extensions of (10) that will improve the efficiency of the estimates and allow them to more accurately reflect the true informational content of the sample. Thus, for example, there is likely to be significant cross country correlation in product specific growth rates as relative prices, globally, follow particular trends. This is incorporated by extending the random effects framework to:

$$(13)\quad \ln(Q_{pct}) = a_p + b_p(c_c + g*t + v_c*t + v_p*t + u_{pc} + e_{pct})$$

---

[6] Of course, in a similar vein one of the product consumption level constants a_p must be dropped to allow the estimation of a full set of country dummies c_c.

where v_p represents correlation in product specific trends. Similarly, the individual products may not be a true random sample, so that there are significant non-income related correlations across groups of products in levels within countries and in growth rates across countries:

$$(14)\quad \ln(Q_{pct}) = a_p + b_p(c_c + g*t + v_c*t + v_p*t + v_G*t + u_{pc} + u_{Gc} + e_{pct})$$

where G denotes a product group (e.g. the housing data in Table I earlier).[7] I have found that there are generally quite significant product group correlations in levels within countries (u_{Gc}), but the estimate of the variance of commonalities in product group growth rates (v_G) is invariably insignificantly different from zero, or the likelihood is maximized when its value is actually set to zero, so I drop this term from the analysis and reported results which follow.

Since the estimated regression coefficients are simple constants and time trends, the random effects of (10), (13) and (14) have little to do with the traditional weighting of "between" and "within" estimators. Instead, large estimated variances of v_c, v_p, u_{pc} and u_{Gc} indicate that there is considerable correlation between the error terms, i.e. that the nominal number of observations considerably overstates the true information in the sample. This naturally produces larger standard errors, but also changes the point estimates as less than the usual OLS weight is placed on observational groups which are highly correlated.[8]

As a final refinement, I should note that efficiency requires that one explicitly take into account the fact that the observations of the left-hand side variable in (14), i.e. product

[7]All shocks are orthogonal to each other. Thus, for example, e_{pc} represents the residual country x product correlation after accounting for country x product group correlation e_{Gc}.

[8]Thus, for example, if there are strong product or country correlations in growth rates then, in calculating the average growth rate g, comparatively less weight is placed on individual observations belonging to products or countries with a large number of total observations. Similarly, if there are strong product or product group correlations in levels then, in calculating the country fixed effects, comparatively less weight is placed on individual observations belonging to products or product groups with a large number of observations in a given country.

consumption levels, are estimated from data in a preliminary first-step. To this end, the covariance matrix for the generalized least squares multivariate normal likelihood should be augmented with the covariance matrix of the first step estimates, i.e. Cov = Σ(RE) + Σ(first step), where the first term is the covariance matrix described by the random effects model (e.g. eqn. (14)) and the second is the covariance matrix of the dependent regression variable, as estimated in the first step. Inserting the first step covariance in this manners functions, in a fashion,[9] like weighted least squares, as less weight is placed on dependent variables which are estimated with more uncertainty. For estimation based upon national averages calculated from individual surveys, in the following section, this is not of great import, as the sample sizes ensure that the standard errors are miniscule. Further on in the paper, however, where micro-level correlations are used to estimate the quasi-income elasticities, this procedure places less weight on products whose association with income is weaker, and hence provide less information about trends and relative levels of real living standards.[10]

Estimation error can seriously bias the inferences one draws about cross-sectional variation in consumption levels. Since the country consumption levels, c_c, of equations such as

[9]In a fashion, because they are not a strict relative weighting. In weighted least squares, observations are weighted by their relative variance, but here the absolute first step variance relative to the model-residual observed variation in the estimated variables is relevant. Thus, when all of the first step errors are small, as in the case discussed below, the implied real variation in the dependent variable dominates and there is, effectively, not much relative weighting. To see this more clearly, consider the case where both the first step covariance and second step model error are diagonal, so that the variance of each observation i is given $\sigma^2(y_i) + \sigma^2(\varepsilon)$, where the first term is the first-step variance of the estimated dependent variable and the second is the model error (estimated in the second step).

[10]I should note that the standard errors of the coefficients estimated in the second-step should also be modified to take into account their dependence on the first-step parameter estimates and their standard errors, i.e a two-step calculation of standard errors. However, given the complexity of some of the likelihoods, the calculations are quite difficult and time consuming. Hence, having checked a few cases and found that this does not substantially alter the results, I leave this particular refinement for a final draft.

(7) and (14) are inevitably estimated with error, their cross-sectional standard deviation is exaggerated by estimation error. One can correct for this by taking these estimates and running a third-step (!) regression of their value on a constant, incorporating the estimation variance into the regression likelihood in the manner described in the preceding paragraph. The estimated regression standard error of this constant-only model would represent the true cross-sectional variation of base year consumption levels. Calculating the theoretically correct standard error for the ratio of second (g) and third (σ) step coefficients, each estimated with dependent variables and embedded covariance matrices estimated in earlier stages, involves equations of grotesque complexity. A simpler approach is to estimate the country consumption levels, in the second-step models described above, as a random effect:

$$(7)' \;\; \ln(Q_{pct}) = u_{pc} + g_p * t + v_{pc} * t + e_{pct}$$

$$(14)' \;\; \ln(Q_{pct}) = a_p + b_p(u_c + g * t + v_c * t + v_p * t + v_G * t + u_{pc} + u_{Gc} + e_{pct})$$

where, as before, u's and v's denote level and growth random effects, respectively, and where, now, u_{pc} substitutes for c_{pc} as a product-specific country level random effect in the single product equation (7) and u_c substitutes for c_c as a cross-product country level random effect in the multi-product equation (14). This random effects specification, combined with the incorporation of the first step covariance matrix in the second step likelihood (as described above), automatically purges the estimate of $\sigma[\ln(C_c^R)]$ of variation due to estimation error. The random effects assumption of independence from the other coefficients, which are simple constants and time trends, is not very demanding and I show, in an appendix available from the author, that the full three step fixed-effects procedure described above produces very similar estimates of the ratio of growth to cross-sectional variation in real living standards.[11] Throughout the paper, I use the random effects specification to estimate directly, in the second step, the cross-sectional variation of real consumption levels.

To summarize, under the assumption of a ln linear income elasticity of sorts, the time series and cross-sectional variation in a measure of the real consumption of a product provides information on the ratio of the growth to the standard deviation of real expenditure. Econometrically, estimation of this ratio can explicitly account, in a variety of ways, for correlation within countries, products and product groups in both levels and growth rates, producing standard errors and coefficient estimates that properly reflect the amount of independent information in the product sample. The use of estimated means as a left hand side variable requires that the first stage covariance matrix be incorporated in the second stage likelihood's description of the covariance matrix. These econometric refinements play quite a significant role later in the paper, when I separately estimate the trend and standard deviation of real expenditure. They are not, however, necessary to establish the central result of this paper, as presented in the next section, i.e. that there is a glaring inconsistency, particularly concentrated in the sub-Saharan economies, between the relative degree of time series and cross-sectional variation in real consumption present in the DHS and the PWT.

[11] However, for the reason described above, I do not calculate the standard error of the three-step estimates presented in that appendix.

IV. Results I: The Gross Inconsistency Between the DHS and PWT

Table II below presents product-level estimates of the ratio of the growth to the standard deviation of real consumption expenditure following equation (7)' described in the previous section. For each entry in the table, the dependent variable is a measure of the mean consumption of a product in a particular survey, i.e. a country x time panel for a given product. For ln rooms per capita and children's ln weight and ln height, this is the sample mean. For the remaining dichotomous variables, coded in the surveys as 0 or 1, the measure used in the table is the logit of the mean value, i.e. $\ln(\bar{x}/(1-\bar{x}))$.[12] I calculate these values at the urban and rural level for each survey and aggregate them using the DHS total urban/rural household weight to produce a national measure.[13] The estimated variance of these first step estimates is adjusted for clustering and then additively incorporated into the likelihood of the second step random effects model, as described earlier above.[14] Although these refinements matter in later sections, in this table they are not crucial, and estimates calculated using simple national means (without urban/rural weighting), with or without adjustment for the first step variance or clustering, and even including more complicated random effects (at the cluster level) estimation of regional

[12]I use the logit as my baseline, rather than simply the sample mean value, because later in the paper I will need to estimate probability equations at the micro level, i.e. deal explicitly with the fact that the variable can only take two values, and the logit is a convenient probability model that is easily extended to random and fixed effects specifications. However, as shown in this section and later in the paper, the results are fairly robust to the choice of functional form.

[13]As for the independent variable "time", since the surveys are executed over a period of months, I code each survey as taking place at the average date at which the households were interviewed, with each month coded as 1/12 of a year. Section II and Appendix I provide further details on the definition and construction of the variables.

[14]To clarify, the first step predicted values are equal to those one would get if one ran, for each urban/rural region in each survey, a linear regression or logit model on a constant. The standard errors of these urban/rural means can then be adjusted for clustering using the usual robust "sandwich" estimator of variance. The national mean for each survey is the weighted average of the urban/rural means, and its variance is the sum of the square of the urban/rural weights times the individual urban/rural variances.

Table II: Product Level Estimates of Growth/Standard Deviation
Dependent variable = urban/rural weighted country means
$y_{pct} = a_p + g_{p\sim A}*t + g_{pA}*t + u_{pc} + v_{pc}*t + e_{pct}$, reporting $g_i/\sigma[u_{pc}]$

Durables	$g_{p\sim A}/\sigma[u_{pc}]$	$g_{pA}/\sigma[u_{pc}]$	Housing	$g_{p\sim A}/\sigma[u_{pc}]$	$g_{pA}/\sigma[u_{pc}]$
Radio	.019 (.015)	.089 (.016)	Electricity	.038 (.006)	.032 (.006)
Television	.040 (.007)	.055 (.008)	Tap Water	.005 (.010)	.008 (.009)
Refrigerator	.034 (.007)	.026 (.006)	Flush Toilet	.052 (.009)	.017 (.008)
Bicycle	.028 (.010)	.048 (.009)	Constructed Floor	.027 (.007)	.016 (.006)
Motorcycle	.031 (.011)	.026 (.010)	ln(Rooms/Capita)	.058 (.012)	-.002 (.007)
Car	.020 (.009)	.021 (.008)			
Telephone	.076 (.017)	.068 (.018)			
Children's Health	$g_{p\sim A}/\sigma[u_{pc}]$	$g_{pA}/\sigma[u_{pc}]$	**Family Economics**	$g_{p\sim A}/\sigma[u_{pc}]$	$g_{pA}/\sigma[u_{pc}]$
ln Weight	.029 (.010)	.031 (.008)	At School (6-14)	.053 (.014)	.069 (.013)
ln Height	.061 (.015)	.039 (.012)	At School (15-24)	.054 (.015)	.043 (.014)
No Diarrhea	.006 (.020)	.058 (.020)	Working (15-24)	.010 (.019)	.019 (.013)
No Fever	.015 (.024)	.100 (.025)	Working (25-49)	.029 (.023)	.033 (.016)
No Cough	.009 (.026)	.075 (.026)	Birth (15-24)	.092 (.020)	.021 (.016)
Alive	.081 (.013)	.037 (.009)	Birth (25-49)	.063 (.008)	.011 (.005)
			Marriage (15-24)	.012 (.009)	.027 (.009)
			Marriage (25-49)	.011 (.006)	.024 (.006)

Note: A = sub-Saharan Africa; ~A = non sub-Saharan Africa. Each product is estimated separately with a pooled global sample and 2000 as the base year for the cross-sectional variance (i.e. t = year - 2000).

means, are all virtually identical. To avoid pointless repetition, these are placed in an appendix, available upon request from the author. For comparison purposes, I estimate separate average growth rates for the sub-Saharan and non sub-Saharan economies divided by the global standard deviation of base year consumption levels.[15] I find, across all 26 products, the average of the growth to standard deviation ratio and its standard error to be .037 (.013) in the non-African economies and .038 (.012) in sub-Saharan Africa.

[15]Obviously it would make no sense to divide African growth by the African variance and try to compare it to the same number for the non-African countries. I should note that, throughout this paper, I treat 2000 as the base year, i.e. t is measured as year minus 2000.

Table III: PWT & UN Based Estimates of the Growth and Standard Deviation of Real Living Standards $y_{ct} = a + g_{\sim A}*t + g_A*t + u_c + v_c*t + e_{ct}$			
	Penn World Tables 6.2 Private Consumption per Capita	Penn World Tables 6.2 Private Consumption per Equivalent Adult	UN National Accounts Private Consumption per Capita
$g_{\sim A}$	.020 (.004)	.017 (.004)	.022 (.004)
g_A	.010 (.004)	.009 (.004)	.009 (.003)
$\sigma[u_c]$	.677 (.065)	.651 (.062)	.709 (.068)
$\sigma[v_c]$	.012 (.003)	.012 (.003)	.011 (.003)
$\sigma[e_{ct}]$	.072 (.008)	.072 (.008)	.080 (.009)
$g_{\sim A}/\sigma[u_c]$	.030 (.007)	.027 (.007)	.030 (.006)
$g_A/\sigma[u_c]$	.014 (.006)	.013 (.006)	.013 (.005)

Notes: Calculated using PWT Laspeyres measures of GDP, with ratio of equivalent adults to population calculated from reported PWT chain measures. PWT chain measures (for GDP) produce identical results. PWT calculates equivalent adults by assigning a weight of .5 to persons under 15. UN measures are in constant market exchange US dollars, with ad hoc PPP adjustments (see text).

Table III runs the same random effects regression used in Table II on the real consumption data of the Penn World Tables. In this case, both the growth and standard deviation of real consumption are identified, as there is no income elasticity that needs to be implicitly estimated, but only their ratio can meaningfully be compared to the DHS results presented up to this point in the paper. I use as my observations the 132 PWT country x year observations corresponding to the countries and dates of my DHS surveys.[16] As the dependent variable I use

[16] I average/weight the PWT data for the years in which each DHS survey takes place (e.g. 2003-2004) based upon the date in which the average household was surveyed. For the countries in my sample, the PWT 6.2 data end, mostly, in 2003 and 2004. For 38 of my DHS surveys the survey date is at or past the last PWT observation for that country (in 10 cases, the survey begins in the last PWT observation year). In these cases, I substitute the last available PWT observation for that country (and its corresponding date). In the case of three countries there are two surveys past the last PWT observation. In those cases, I drop one observation for each country. In sum, the PWT sample consists of 132 country x year observations, with 25 representing data before the corresponding DHS date (22 of these being within two years). The UN data extend to 2006, and hence can match all of the 135 survey x year observations of my DHS data.

per capita and per equivalent adult measures of real private consumption. As the reader can see, the PWT ratio of growth to standard deviation for these measures is .030 for the non-African economies and .014 for sub-Saharan Africa. For sub-Saharan Africa this is close to 1/3 the average present at the product level in the DHS data where, as shown in Table II, 23 of the 26 products show ratios higher than .014. It is immediately apparent that there is a vast discrepancy, concentrated in sub-Saharan Africa, between the degree of growth to cross sectional variation present in the PWT and that present in the DHS. As a cross-check, Table III runs the same analysis on the UN National Accounts Main Aggregates Database measures of GDP in constant US dollars. The UN growth and standard deviation measures are extremely close to PWT. This is not surprising as, given the benchmark levels of expenditure, PWT extrapolates international dataset measures of growth by GDP component, while the UN database, despite being nominally at market exchange rates, makes ad hoc PPP adjustments to levels.[17]

The results reported above understate the true magnitude of the discrepancy between PWT and the DHS, as the DHS product level regressions understate the ratio of growth to standard deviation of real consumption expenditure by including product level cross-sectional variation in the denominator, as explained earlier in equation (9)'. To this end, Table IV presents DHS random effects regressions run across all products, together, following equation (14)' presented earlier.[18] With the product level shocks purged from the country level standard

[17] In the case of economies with volatile price levels and exchange rates, an adjustment is made using relative domestic/US inflation rates back to "the year closest to the year in question with a realistic GDP per capita US dollar figure" (http://unstats.un.org/unsd/snaama/formulas.asp).

[18] To get things started, the first column of the upper panel is (14)' without any random effect other than that representing country levels (u_c). This regression is inefficient, as the large country x product specific error is included in the general error term, but, relative to the results presented in Table II, accomplishes the basic task of averaging out, i.e. eliminating the country x product specific shocks in the implicit calculation of the cross-sectional variation. Additional random effects are added as one moves to the right, as indicated by the coefficients reported in the table.

Table IV: DHS Estimates of Growth/Standard Deviation Dependent variable = urban/rural weighted country means Panel data: product x country x time observations $y_{pct} = a_p + b_p*(u_c + g_{\sim A}*t + g_A*t + v_p*t + v_c*t + u_{pc} + u_{Gc} + e_{pct})$				
	All Products			
$g_{\sim A}/\sigma[u_c]$	.055 (.009)	.053 (.006)	.055 (.009)	.056 (.009)
$g_A/\sigma[u_c]$	.052 (.009)	.053 (.006)	.054 (.009)	.054 (.009)
$\sigma[u_{pc}]/\sigma[u_c]$		1.20 (.124)	1.20 (.122)	1.13 (.121)
$\sigma[v_p]/\sigma[u_c]$			.021 (.004)	.022 (.005)
$\sigma[v_c]/\sigma[u_c]$			.020 (.004)	.021 (.004)
$\sigma[u_{Gc}]/\sigma[u_c]$				.440 (.072)
$\sigma[e_{ct}]/\sigma[u_c]$	1.19 (.118)	.434 (.045)	.396 (.041)	.398 (.043)
	Consumer Durables	Housing	Children's Health	Family Economics
$g_{\sim A}/\sigma[u_c]$	.049 (.012)	.043 (.011)	.049 (.015)	.075 (.017)
$g_A/\sigma[u_c]$	.066 (.012)	.022 (.010)	.068 (.015)	.063 (.015)
$\sigma[u_{pc}]/\sigma[u_c]$	.969 (.116)	.840 (.106)	.876 (.112)	1.82 (.278)
$\sigma[v_p]/\sigma[u_c]$	.020 (.007)	.016 (.006)	.013 (.008)	.025 (.009)
$\sigma[v_c]/\sigma[u_c]$	.026 (.006)	.024 (.005)	.041 (.010)	.020 (.009)
$\sigma[e_{ct}]/\sigma[u_c]$	.255 (.030)	.206 (.025)	.547 (.067)	.598 (.088)

Notes: All equations include a random effect at the country level (u_c) normalized to have a standard deviation of 1. Subscripts indicate the categories within which the shocks operate and across which the shocks are uncorrelated. The product group regressions cannot be executed with group level random effects (Gc), as these would be colinear with the other terms in the regression. All equations include the first step coefficient covariance matrix as an additive part of the second step covariance likelihood. See section III for further details.

deviation by the use of a panel, the average growth to standard deviation across all 26 products is now seen to be about .055 around the world, i.e. almost 2 times that present in PWT for non sub-Saharan countries and roughly 4 times PWT's figure for sub-Saharan Africa. As shown in the table, this result is not due to the dominant influence of a particular product group as all product groups indicate a growth to standard deviation ratio vastly above that in PWT. The lowest group, housing in sub-Saharan Africa, brought down by the outlier of rooms per capita (Table

II), still shows more than 1.5 times the movement to variation present in PWT. While the growth and level country and product random effects have some influence on the coefficient estimates and their standard errors, they are not essential to the overall results, as can be seen by comparing the different columns of the table.[19]

The results reported above use a particular functional form, the logit, to evaluate the dichotomous variables that form most of my product sample. The reader might be concerned that the ln odds ratio this functional form produces transforms variation in mean values to variation in measured consumptions in such a way as to generate the results reported above. To explore this, I consider, as alternatives, the probit, weibull, cauchy and linear probability models. The predicted probabilities and estimated dependent variables associated with each functional form are presented in Table V. The probit has slightly thinner tails than the logit, the cauchy has dramatically thicker tails, the weibull is not even symmetric[20], and the linear model, of course, is simply a linear regression. The coefficient of variation, skewness and kurtosis of the 2912 product x country x survey consumption level estimates for dichotomous variables produced by these functional forms differ substantially. However, as shown in the bottom row of the Table, when the random effects regression of the upper right hand corner of Table IV earlier is estimated, the different functional forms yield very similar results. The relative time series and cross sectional variation present in the DHS consumption data is quite insensitive to the functional form used to convert means to consumption indices.

[19]These random effects adjustments have, as one might expect, a much greater influence on the product group estimates, where absent the implicit RE reweighting, individual product or country observations have a much greater influence. Thus, for example, the sub-Saharan growth to standard deviation ratio for housing, absent any random effects at all, is .009.

[20]Symmetry being that 1-F(y) = F(-y).

Table V: Alternative Functional Forms					
	Logit	Probit	Weibull	Cauchy	Linear
Predicted Probabilities	$\dfrac{e^{y}}{1+e^{y}}$	$\dfrac{1}{\sqrt{2\pi}}\int_{-\infty}^{y} e^{-v^2/2}dv$	$e^{-e^{y}}$	$\dfrac{1}{\pi}\tan^{-1}(y)+\dfrac{1}{2}$	N.A.
Dependent Variable	$\ln\left[\dfrac{\bar{x}}{1-\bar{x}}\right]$	$F^{-1}(\bar{x})$	$\ln(-\ln(\bar{x}))$	$\tan[\pi(\bar{x}-1/2)]$	$\bar{x}$
Coefficient of Variation	4.87	5.60	3.16	8.58	0.65
Skewness	-0.57	-0.38	-0.47	-18.23	-0.08
Kurtosis	3.01	2.36	3.35	443.40	1.63
$g_{\sim A}/\sigma[u_c]$	.056 (.009)	.057 (.009)	.062 (.009)	.051 (.009)	.060 (.009)
$g_{A}/\sigma[u_c]$	.054 (.009)	.053 (.009)	.048 (.008)	.051 (.008)	.050 (.008)

Notes: For the dichotomous 0-1 variables, let y denote the probability index, so that the first row equals the probability of a 1. Inverting this function produces the second stage dependent variable (y) as a function of the mean sample outcome, as described in the second row. N.A. - not applicable, the linear model is not a predicted probability (between 0 and 1), but simply a linear regression. F^{-1}, inverse cumulative standard normal. Coefficient of Variation - average second central moment (of the 2912 product x country x survey probability indices for dichotomous variables) divided by the absolute value of the mean; Skewness & Kurtosis - average third and fourth central moments divided by the standard deviation raised to the third and fourth power, respectively. $g_i/\sigma[u_c]$ = estimated value using specification of column (4) in Table IV across all products (including the non-dichotomous ln rooms, weight and height entered, in all specifications, as sample means).

Figure I below illustrates, graphically, the data underlying the results reported above and later in this paper. For each product x survey combination, I graph the country demeaned values of the product consumption levels[21] against the country demeaned values of the survey year.[22]

[21]Again, for the ln variables (rooms, height and weight) this is simply the urban/rural weighted regional average, whereas for the dichotomous variables it is the urban/rural weighted logit (or folded ln) of the regional mean, $\ln(\bar{x}/(1-\bar{x}))$.

[22]In each figure, I drop the (usually 14) countries for which I have only one survey observation on the product in question, as they play no role in estimating trends. The data of these surveys contribute to the

To benchmark product consumption levels to overall consumption expenditure, I scale each product measure so that the standard deviation of the country means equals the estimated PWT standard deviation of consumption per equivalent adult (.651 in Table III above), so that the vertical axis in the figures represents PWT-comparable consumption growth.[23] Examining the figure, it is immediately apparent that, across virtually all products, there is simply "too much" movement relative to cross-sectional variation, particularly for the African countries. Based upon PWT and UN growth rates, a country (demeaned) year value of 5 should be associated with observations below .05 for Africa, i.e. a negligible movement on the vertical scale of the graph. This is clearly not the case, with most products showing robust growth. Either Africa is growing much faster than indicated by standard international sources, or the cross-sectional standard deviation is much lower than indicated by PWT (so that the vertical movement should be scaled down). In sections V and VI below I extend my methodology to allow the separate estimation of both the growth rate and the standard deviation of consumption expenditure and, ultimately, conclude that the discrepancy lies in growth. As will be seen, individuals with a strong prior that the return to education is very low could, methodologically, scale my results and eliminate the growth discrepancy, but only at the expense of suggesting that the cross-sectional variation in living standards is much smaller than indicated by PWT.

Comparing the left and right hand panels for each product in Figure I, it is also apparent that it is hard to find evidence that Africa is growing slower than other developing countries. There are a few products in which African living standards appear to be improving more slowly,

---

estimation of the cross-sectional standard deviation of consumption, here and later, and to the micro-data estimation of quasi-income elasticities and demographic effects later in the paper.

[23]Thus, if y_{it} is the country demeaned consumption measure, y_i the country mean consumption measure, and σ[PWT] the PWT standard deviation of ln real consumption, I scale each y_{it} by σ[PWT]/ $\sigma[y_i]$.

Figure I: Product Level Consumption Growth
(Cross-Sectional Standard Deviation Normalized to PWT Levels)

Durables & Housing Conditions

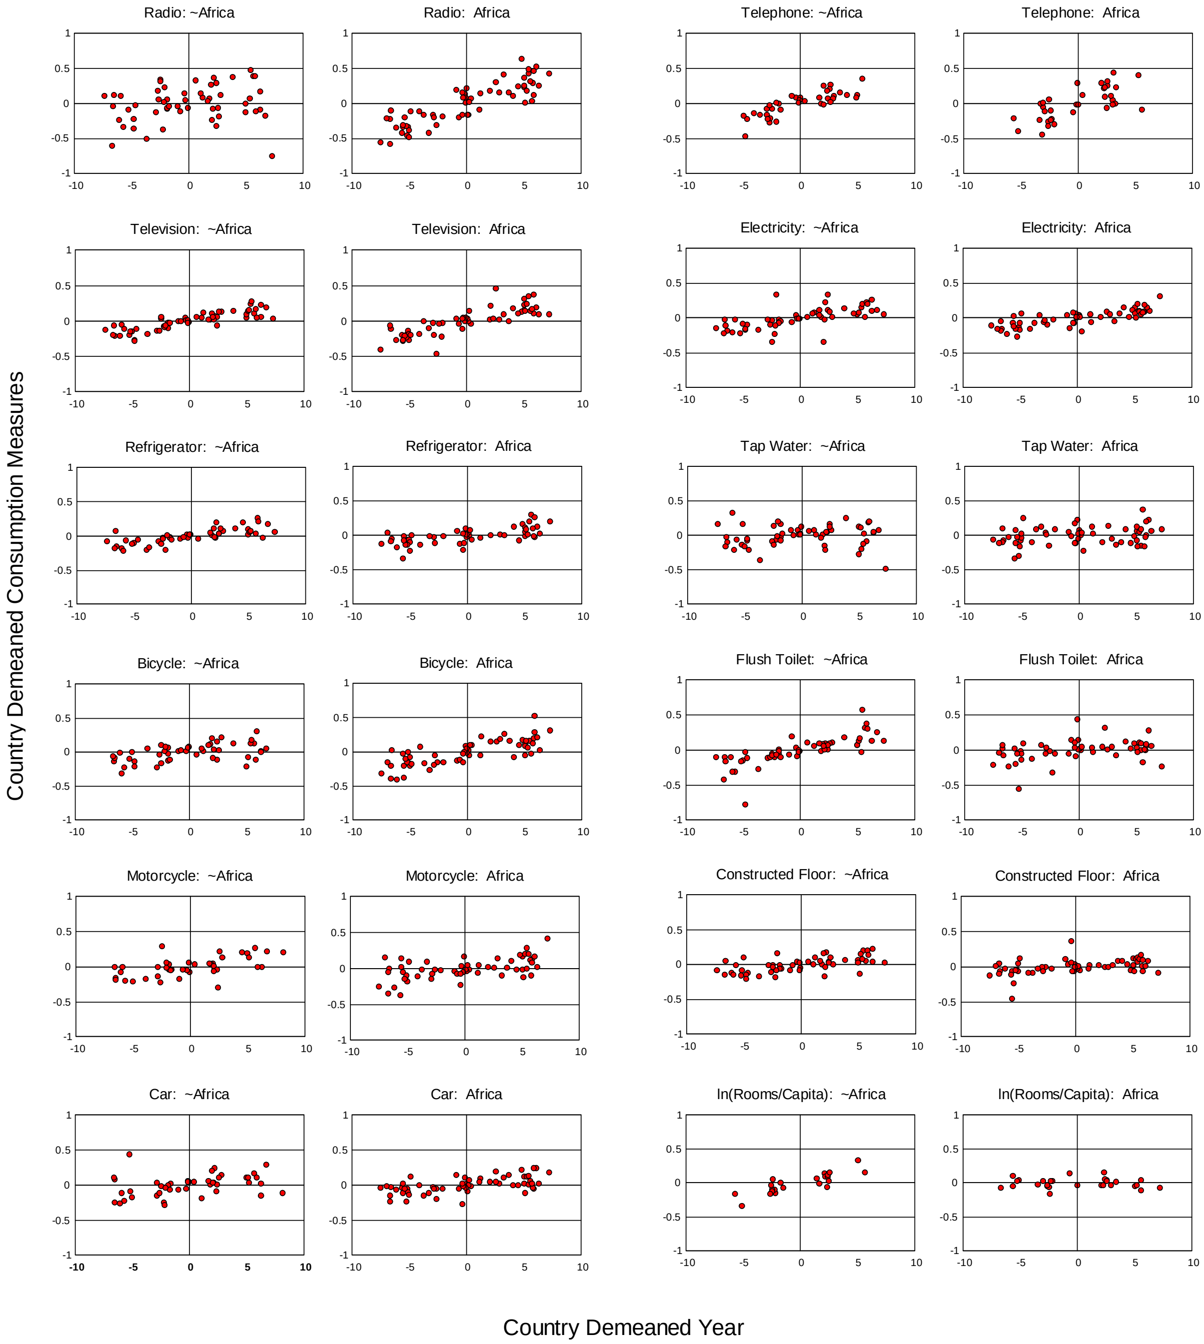

Figure I: Continued

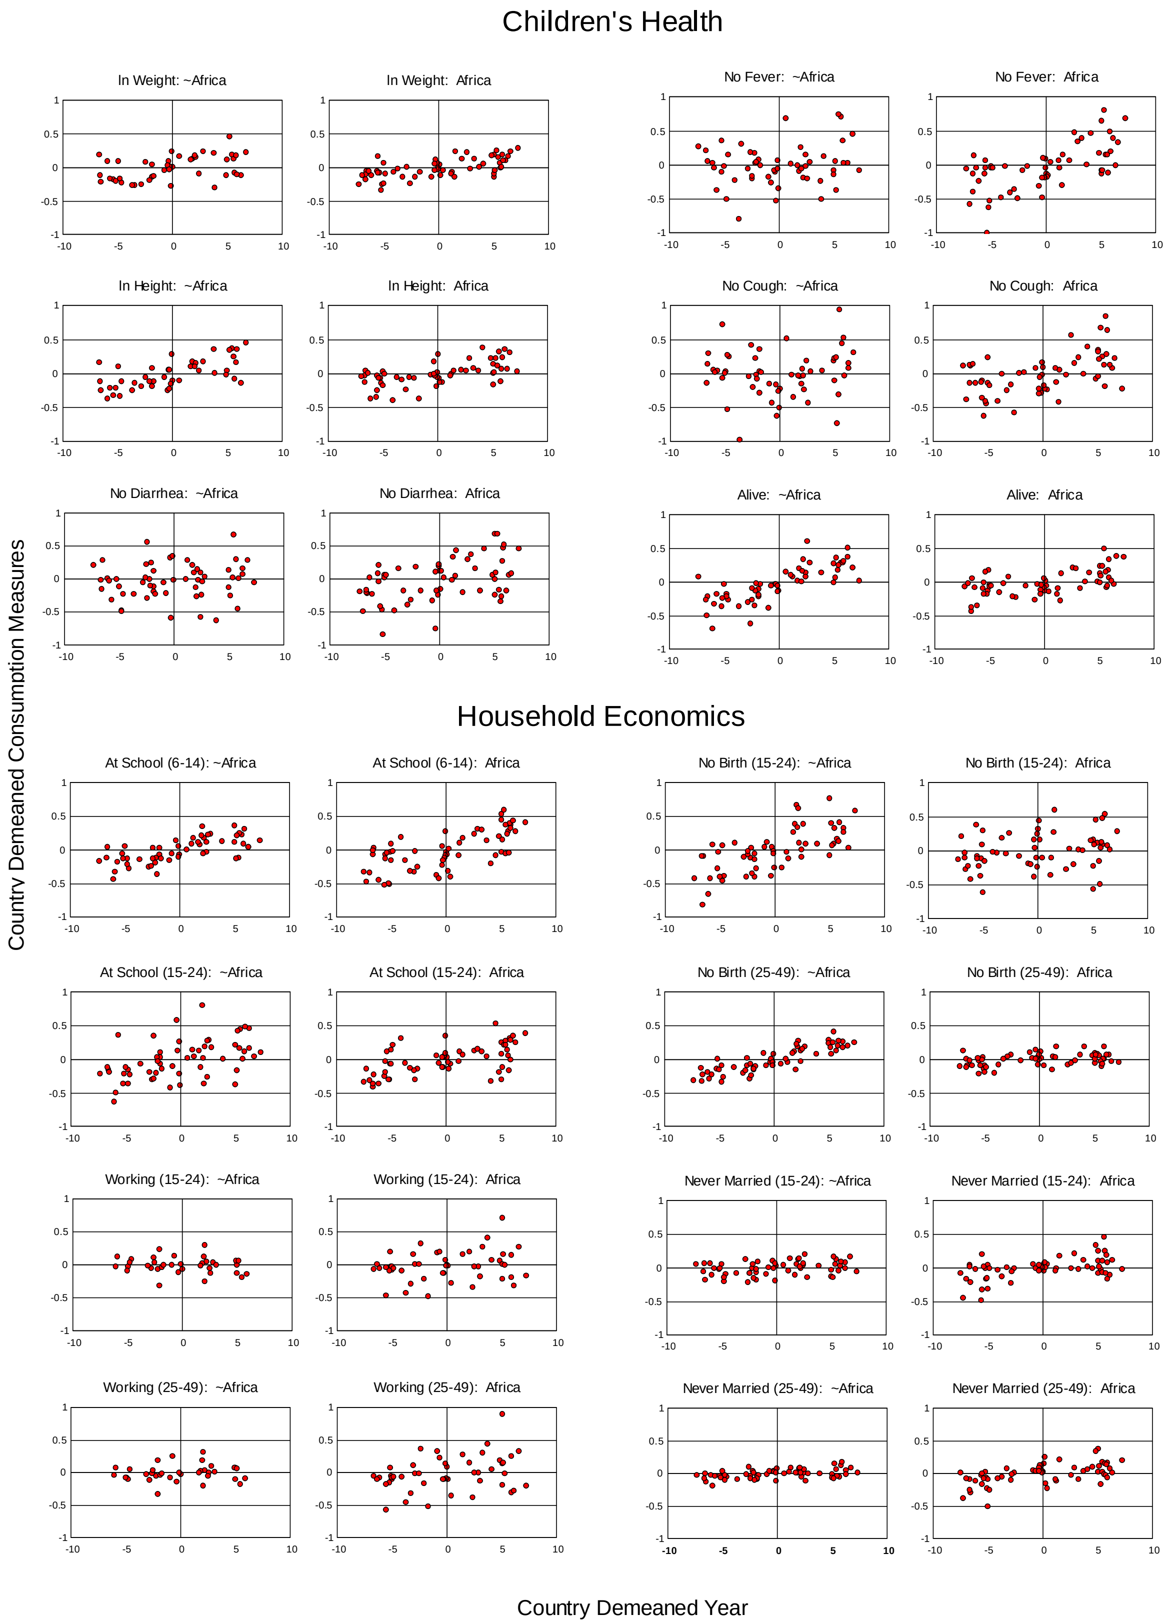

as well as a few products in which they appear to be improving more rapidly, but one would be hard pressed to find any systematic difference between the African and non-African panels. Clearly, no amount of uniform scaling can eliminate this result. It is possible, however, that there is heterogeneity in demand patterns, so that the vertical movements in the African graphs represent smaller underlying movements in real expenditure, as would be the case if African income elasticities for these products were systematically larger. I explore the issue of heterogeneity in sections VII and VIII later in the paper. While demand patterns do vary substantially across countries, they do not vary in a way that is systematically correlated with these results, so that I find that, after allowing for heterogeneity in demand, African growth remains comparable to that of the non-African developing economies.

V. Methods II: Incorporating Micro Correlations

Let the real demand by household h for product p in region r in period t be described by the equation:

$$(15)\quad \ln(Q_{hprt}) = \alpha_p + \eta_p \ln(C^R_{hrt}) + \vec{\beta}'_p \vec{X}_{hrt} + \eta_p \varepsilon^{\vec{P}}_{prt} + \varepsilon_{hprt}$$

which is merely (4) earlier specified at the micro household level and augmented with a vector of household demographic variables X_{hrt}, like household size and age composition, that shift demand through the coefficients β_p. The error term is now made up of two components, the influence of relative prices, whose effect, as explained earlier, is proportional to the quasi-income elasticity of demand, and a mean zero idiosyncratic household preference shock. C^R_{hrt}, household real consumption expenditure per adult, can reasonably be thought of as being proportional to permanent income per adult, which in turn is related to educational attainment:

$$(16)\quad \ln(C^{R}_{hrt}) = \alpha_{rt} + \ln(Y^{R}_{hrt}) \qquad \ln(Y^{R}_{hrt}) = \ln(Y^{R\sim E}_{rt}) + R_E E_{hrt}$$

where E_{hrt} is the average years of educational attainment of adult household members, R_E is the return to a year of education, and $\ln(Y^{R\sim E}_{rt})$ is education adjusted ln regional real income at time t. It follows that average regional ln household consumption expenditure at time t is given by:

$$(17)\quad \ln(C^{R}_{rt}) = \ln(C^{R\sim E}_{rt}) + R_E \overline{E}_{rt}$$

where $\overline{E}_{rt}$ is mean household educational attainment and $\ln(C^{R\sim E}_{rt})$ is education adjusted ln regional real expenditure per adult.[24]

For each product, combine a number of household level surveys to estimate the equation

$$(18)\quad \ln(Q_{hprt}) = a_{prt} + b_p R_E E_{hrt} + \vec{c}\,'_p X_{hrt} + e_{hprt}$$

where the a_{prt}s are a complete set of product specific region x time (or, equivalently, survey) dummies. Regions, r, can be defined at any level that allows consistent aggregation across time, and, in my case, will consist of the urban and rural areas of each country. Clearly, the estimates of b_p and $\vec{c}_p$, identified off of cross sectional variation within surveys, will be unbiased, but the dummies will capture all common product x region x time components:

$$(19)\quad \hat{b}_p = \eta_p \qquad \hat{\vec{c}}_p = \vec{\beta}_p \qquad \hat{a}_{prt} = \alpha_p + \eta_p \ln(C^{R\sim E}_{rt}) + \eta_p \varepsilon^{\bar{P}}_{prt}$$

While the unconditional expectation of $\varepsilon^{\bar{P}}_{prt}$, the influence of relative prices, is zero, it takes on particular values within any particular product x region x time grouping and ends up being incorporated into the constant term.

Finally, construct measures of ln real regional consumption expenditure at time t, as implied by the consumption of a particular product, as the sum of the product x region x time

---

[24] Clearly, there is no assumption that savings rates are equal across regions (as α_{rt} is incorporated in C^R), but there is the implicit assumption that savings rates out of permanent income do not vary by educational attainment. If they do, R_E has to be changed to reflect this.

first stage dummy divided by the income elasticity of demand, plus the ln real consumption attributable to the separately estimated average regional educational attainment:

$$(20)\quad \ln(\hat{C}^{R}_{prt}) = \frac{\hat{a}_{prt}}{\hat{b}_{p}} + R_{E}\hat{\bar{E}}_{rt}$$

Weighted using the regional household population shares, these measures produce a panel dataset of country consumption expenditures, as measured by different product equations:

$$(21)\quad \ln(\hat{C}^{R}_{pct}) = \sum_{i \in c(r)} S_{i} \ln(\hat{C}^{R}_{pit})$$

where c(r) is the set of regions in country c (in my analysis, the urban and rural areas). These measures can then be projected, in a random effects panel regression on product dummies, an average international growth rate, and random effects capturing the variation in international growth rates and cross correlations in the error term brought about by levels and trends in relative prices:

$$(22)\quad \ln(\hat{C}^{R}_{pct}) = a_{p} + g*t + u_{c} + v_{c}*t + v_{p}*t + v_{G}*t + u_{pc} + u_{Gc} + e_{pct}$$

The reader will recognize this as no more than equation (14)' earlier in section III divided by the quasi-income elasticity of demand, which is now estimated in the first stage regressions which produce the estimates of the dependent variable.[25] Asymptotically

$$(23)\quad \hat{a}_{p} = \alpha_{p} / \eta_{p}\ , \quad \hat{g} = g\ , \quad \hat{\sigma}[u_{c}] = \sigma[\ln(C^{R}_{c})]$$

where g is the average international growth rate of real consumption and $\sigma[\ln(C^{R}_{c})]$ is the standard deviation of country real consumption expenditure in the base year (2000), as described by (5) earlier in section III.

[25] As before, I find that the estimate of the variance of common product group growth rates (v_G) is invariably insignificantly different from zero, so I remove this effect from the final estimating equations reported below.

(a) Details and Extensions

A quick examination of equations (18) and (20) above reveals the obvious fact that the estimated quasi-income elasticity, $\hat{b}_p$, is inversely related to the educational profile of household income, R_E, so that $\ln(\hat{C}^R_{prt})$, real regional x time consumption as defined by product p, is linear in R_E. It follows that all of the equations, up through and including the random effects panel regression (22), used to estimate the growth and standard deviation of living standards, can be estimated in terms of years of education (i.e. setting R_E temporarily equal to 1) and the results then multiplied by one's estimate of R_E to arrive at income equivalent measures of growth or variation. I highlight this fact so that the reader can see that any disagreement with my estimate of R_E from the DHS data (.116, as shown later below) can be resolved by simply scaling the estimates and standard errors proportionately to one's preferred number.[26]

As in section III, in estimating (22) I add the first step covariance matrix for the dependent variable, $\ln(\hat{C}^R_{pct})$, to the random effects covariance matrix in calculating the covariance matrix for the GLS maximum likelihood. This is not only justifiable econometrically, as a means of improving the efficiency of the estimates, but also, in more pedestrian terms, makes the procedure somewhat "idiot proof". Products where the estimated relationship between relative incomes (education) and demand is statistically weak will have

[26]For this reason, I treat R_E as "known" and do not incorporate the standard error of its estimate in calculating the standard error of the other coefficients reported below. Doing so, however, has a negligible effect on the reported standard errors. Since my estimate of R_E is calculated independently of the other coefficient estimates it is easily shown that the actual variance of each coefficient B is given by:

$$Var(\hat{B}) + [(\hat{B})^2 + Var(\hat{B})][Var(\hat{R}_E)/(\hat{R}_E)^2]$$

where $Var(X)$ is the variance of the estimate of X, as reported in the tables below. As $Var(\hat{R}_E)/(\hat{R}_E)^2 = .0017$, and the coefficients of interest on growth and the standard deviation of living standards are on the order of about 5 to 10 times their standard errors, adjustments along these lines would multiply the reported standard errors by a factor of between 1.02 and 1.08.

very large first step standard errors. In this regard, putting aside the random effects covariance terms, the GLS likelihood will function like weighted least squares, discounting the variation in those observations. Consequently, randomly picked real choice variables, such as the household's favourite colour, whose association with real incomes is dubious, will have little effect on the estimates of real relative consumption expenditure.[27]

The inclusion of the first-step covariance matrix for $\ln(\hat{C}^{R}_{pct})$ in the second step GLS likelihood raises an important technical stumbling block. As shown in (20) above, $\ln(\hat{C}^{R}_{prt})$, which is used to construct $\ln(\hat{C}^{R}_{pct})$, is computed as the ratio of normally distributed variables. In calculating the distribution of non-linear functions of normal variables, it is customary to make use of the "delta method", an application of the central limit theorem. However, even the central limit theorem has its limits. As the probability mass around zero of the random variable in the denominator increases, the central limit theorem breaks down, the most notable example of which is the well known result that the ratio of two independent standard normal variables follows a cauchy distribution, which doesn't even have any moments. To the degree that the denominator in (20), the quasi-income elasticity, differs from zero, this is not a problem as, asymptotically, the probability mass of the estimated coefficients around zero goes to zero. However, for finite samples, or in the case of poorly chosen variables whose correlation with real incomes is spurious, the probability mass around zero can be large enough to make the delta method calculation of the first step covariance matrix utterly inaccurate. I handle this problem by using Monte Carlo techniques to estimate the first step covariance matrix of $\ln(\hat{C}^{R}_{prt})$.[28] For

[27]However, the estimates are not protected from bias introduced by the use of real variables, such as race, which are strongly correlated with education and incomes but not subject to individual choice.

[28]To be clear, I accept the standard maximum likelihood estimates of the first step coefficients and their covariances, as these are based upon 100,000s of observations and do not involve ratios of normals. But, in calculating the distribution of (20), I generate a million draws from the estimated joint distribution

almost every single one of my results, the two or three minimal exceptions being noted in footnotes later, the effect is negligible, i.e. the results are virtually identical to those arrived at using the delta method covariance matrix.[29]

Household survey data are collected in "clusters", i.e. groups of households at particular survey locations. This suggests the likelihood of correlation across the error terms for households in the same cluster which makes the first step covariance matrix inaccurate and the coefficient estimates inefficient or, worse, biased. I address this problem in three ways. First, as a baseline, I ignore the clustering in the first step estimation procedure, but calculate more accurate "clustered" first step standard errors using the usual sandwich covariance estimator. Second, I formally estimate first step cluster-level random effects regression and discrete choice logit models. Third, to allow for the possibility that the cluster errors are correlated with the independent variables, I estimate cluster-level fixed effects models.[30]

---

of the a_{prt}s and b_p in each product equation and then calculate the resulting mean and variance of the ratios, to which I then add the covariance matrix of R_E times the estimated mean educational attainment by region.

[29]This is not to say that this problem is never going to be relevant. In considering alternative examples, I find that when the t-statistic of the denominator falls to about 2 the resulting estimate of variance is 1000s of times larger than implied by the delta method. I should note that in statistical theory there has been little progress on distributions of this type, beyond extending the Cauchy result to noting that the distribution of the ratio of two correlated normals produces a distribution that takes close to a page of text to write down. Thus, analytically calculating the finite sample multivariate distribution of (20) is not an option.

I should note that this problem does not invalidate the statement earlier above that the incorporation of the first step covariance matrix in the second step likelihood protects the estimates from "idiot" variables which are not actually correlated with real incomes. As the probability mass of the estimated quasi-income elasticity around zero grows, the estimated Monte Carlo variance explodes to infinity (much faster than indicated by the central limit approximation) and, even though strictly speaking the resulting distribution is no longer normal, the inclusion of the exploding covariance matrix in the GLS normal likelihood places a vanishing weight on observations in that product group.

[30]For the dichotomous variables, I use Butler & Moffitt's (1982) random effects specification, modelling the random effect as normally distributed and using Gauss-Hermite quadrature to integrate the cluster joint logit probability, while for fixed effects I use Chamberlain's (1980) conditional logit likelihood, implicitly differencing out the cluster fixed effects (without actually estimating them) by evaluating the likelihood of a particular cluster outcome conditional on overall cluster characteristics.

Each successive cluster level model I apply is, statistically speaking, found to be superior to the one before, i.e. no random effects are rejected in favour of significant random effects and random effects (and their assumption of independence from the other regressors) are rejected in favour of fixed effects. The large correlation of errors within clusters in the random effects specification places greater weight on "within" cluster variation in educational attainment and consumption levels, which the fixed effects specification completes by looking only within clusters. This produces, empirically, smaller estimates of the quasi-income elasticity of demand and, by extension, greater estimates of the cross sectional and time series variation in living standards. However, it is not clear these estimates are an improvement on those found ignoring cluster level correlations. First, as one tunnels down to the cluster level, the noise to signal ratio in measures of household educational attainment rises, biasing the coefficients towards zero. Thus, it is not clear whether the smaller estimates of quasi-income elasticities of demand are more accurate representations of reality. Second, much of the correlation within clusters in consumption represents, in fact, the outcome of demand (for communal infrastructure) that is implicitly paid for through the cost of housing and land. To this extent, one would clearly want to identify the quasi-elasticity of demand using between cluster, rather than within cluster, variation. For these reasons, I treat estimates without adjustment for cluster random or fixed effects as my baseline,[31] reporting the others as variations for the reader's consideration.

As for both logit and regression the regional dummies cannot be directly estimated with cluster fixed effects, I employ a two-step procedure: first, estimating the income elasticity and demographic coefficients using cluster fixed effects, and then using these estimated coefficients as an offset in a cluster random effects specification where I calculate the regional product dummies. This is justified on the obvious grounds that the cluster errors, within regions, are orthogonal to the regional means. The standard errors of the regional dummies and their covariance with the estimated income elasticity are adjusted for the two-step procedure.

[31]In this case, the correlation within clusters influences the standard errors, which recognize that it diminishes the effective size of the sample, but is not allowed to influence the coefficient estimates.

VI: Results II: The Standard Deviation & Growth Rate of Living Standards

(a) The Return to Human Capital

As a preliminary, I use DHS data on individual earnings from work to calculate the return to education. I focus on individuals 25 or older, whose education can be taken as completed, reporting earnings from working for others (i.e. not for family or self). I find earnings data of this sort for adult women in 26 DHS surveys in 14 sub-Saharan African and 10 other countries, and for adult men in a sub-sample of 16 of these surveys in 11 sub-Saharan countries and 5 other countries (see Appendix I). I run the typical Mincerian regression of ln wages on educational attainment, age, sex and regional controls.[32]

As shown in Table VI, the OLS estimate of the return to human capital is somewhat sensitive to the number and level of regional controls. While column (1) includes the most basic controls, a dummy variable for the nominal level of wages in each survey, column (2) includes survey x rural/urban controls. Doubling the number of locational controls in this fashion lowers the return to a year of education from 11.5 to 10.8 percent. Adding random effects at the cluster level (column 3) lowers the marginal return further, while fixed effects at the cluster level (column 4) bring it down to 9.5 percent. These results can be rationalized by arguing that rich people tend to live together in rich places, i.e. regions and locales (such as urban centres) which provide higher earnings for any given level of education. As more detailed locational controls are introduced, the return to education is increasingly identified from within locale differences in educational attainment and incomes, rather than cross regional income differences. However, it

[32]Strictly speaking, in my model, R_E refers to the education profile of ln permanent income. Under the assumption that labour income provides the best measure of this, and to produce a "global" Mincerian regression that readers will find comparable to other sources, I restrict the analysis to earnings from working for others. As noted in a later footnote, broader measures of income produce a somewhat higher estimate of R_E.

Table VI: Ln Wage Regressions					
	(1) survey dummies	(2) survey & rural/urban dummies	(3) cluster random effects	(4) cluster fixed effects	(5) cluster fixed effects (IV)
educ	11.47 (.154)	10.78 (.148)	10.40 (.127)	9.46 (.156)	11.60 (.477)
age	4.73 (.692)	4.68 (.675)	4.93 (.614)	4.81 (.701)	4.58 (.819)
age2	-0.05 (.009)	-0.05 (.009)	-0.05 (.008)	-0.05 (.009)	-0.04 (.011)
sex	-35.04 (1.95)	-35.95 (1.91)	-36.52 (1.52)	-36.65 (1.70)	-39.56 (1.99)
N	22996	22996	22996	22996	18418

Notes: Dependent variable = 100*ln annualized wage income of individuals 25 or older working for others, so the coefficients can be read as derivatives expressed in percent. Educational attainment and age measured in years, while sex = 1 if female.

is also important to note that more detailed locational controls increase the noise to signal ratio in educational attainment, biasing the coefficient towards zero. This is particularly relevant for the estimates with cluster fixed effects, as these dummies account for 58 percent of the residual (orthogonal to the individual controls) variation in individual educational attainment.

Column (5) of Table VI controls for the measurement error in individual educational attainment by instrumenting it with the mean educational attainment of other adult members of the same household, as well as their mean age, age2 and sex.[33] As shown, when instrumented, the estimated return on human capital jumps to 11.6 percent. When compared with the coefficient for column (4), this suggests that measurement error accounts for about .19 of the residual variation in individual educational attainment in that specification.[34] This would imply a

[33]The absolute values of the t-statistics of these four variables in the first stage regression are 90.5, 3.4, 6.2, and 13.4, respectively.

[34]As is well know, attenuation bias when one variable is measured with error is equal to S/(S+N), where S is the orthogonal (to other regressors) signal variation and N is the noise variation. In the sentence which follows, I multiply 1 minus this ratio times the residual variance estimate (adjusted for degrees of freedom) of the regression of E on cluster fixed effects and the other individual level controls. I should note that the sample of column (5) is smaller, because many individuals live in households

measurement standard error of about 1.6, i.e. that about 36 percent of the wage reporting sample, with mean educational attainment of 9.5 years,[35] over or understate their educational attainment by 1.6 years or more. This is large, but by no means implausible. Adjusting the coefficient of column (2) by this estimate of measurement error produces a point estimate of a "noise adjusted" return[36] to education of 12.5 percent in that column. When compared with column (5)'s point estimate, this indicates that while measurement error is a concern, there is also substantial correlation, below the urban/rural level, between individual's incomes and the education-adjusted income level of the locales they live in.

In what follows, I will take 11.6 percent as my "known" estimate of R_E. Psacharopolous

---

without other adults, and hence cannot be instrumented. The coefficient of column (4) using the sample of column (5) (which is what I use to calculate the signal to noise ratio) is 9.40.

[35]The wage reporting sample is considerably better educated than the average for the adult men and women in the male & female survey modules from which the data come (5.0 years). Most of this selection has to do with working for others, rather than working per se. Thus, the average educational attainment of adults who report they are working is 5.3 years, while the average educational attainment of adults who report earnings data, whether working for themselves or others, is 6.6 years. If I rerun the analysis of Table VI using all individuals reporting earnings from work (including, presumably, capital income) I get education coefficients of 11.1 and 9.2 for the specifications of columns (2) and (4), and 13.6 for the IV regression of column (5) (with an implied measurement standard error of 1.9). Thus, a broader sample with a broader measure of income produces a higher estimate of R_E and, hence, implies a greater discrepancy between the DHS and international measures of growth.

It would be nice to implement selectivity bias adjustments to correct for selection into employment. Conceptually, these are difficult to implement meaningfully in a Beckerian framework in which family economics is part of household demand, so that traditional labour market selection instruments like marital status and pregnancy are seen to be correlated with the relative productivity of the individual in the household and in the market. Nevertheless, to do what is possible (with the DHS data), I have proceeded blindly, augmenting the earnings equation with separate male and female selection equations, including variables such as marital status, current pregnancy (of a woman or a man's partner), and births in the past year, estimating (in an MLE framework) separate correlations between the disturbance terms for these male/female equations and the earnings equation. I consider two possible cases: (1) selection into participation/employment alone, whether working for others or not (with the wage equation focusing only on those working for others, this being taken as random conditional on employment); (2) selection into reporting wage earnings working for others. Working on the specification of column (2), which is the easiest to implement in this framework, I find that the coefficient falls from 10.8 to 10.7 in the first case and rises to 12.1 in the second.

[36]Arrived at by calculating the estimated residual variation of E when regressed on the other controls of column (2), subtracting the measurement error variance noted above, and adjusting the estimated coefficient for the implied attenuation bias.

(1994) in his oft cited survey of Mincerian regressions, finds an average marginal return of 13.4 percent in 7 studies of sub-Saharan Africa and 12.4 in 19 studies of Latin America and the Caribbean, regions which, together, make up 3/4 of the countries in my sample. Thus, the number I use is not particularly large or out of keeping with the existing literature. Readers who have strong alternative priors can simply scale all of the growth rates and standard deviations of real expenditure reported below by the ratio of their preferred number to 11.6. However, as will be seen, it would take an enormous reduction in the estimated return to education, to about 3 percent, to bring the DHS-implied African growth figures in line with international estimates. Moreover, such a reduction would simply shift the DHS-PWT discrepancy from growth to cross sectional variation, producing a new puzzle, as the DHS data would then imply about 1/4 of the cross-sectional variation in levels of expenditure reported in PWT.

(b) First Step Estimates

Table VII below reports the coefficients on household mean years of adult educational attainment in product by product demand equations, estimated with country x survey x urban/rural dummies and household and individual demographic controls, as listed in the notes to the table. With the exception of ln weight, height and rooms per capita, the dependent variable in each row is a 0/1 dichotomous variable and the reported figures represent the coefficients in a logit discrete choice model. The second and third columns of the table run the baseline specification with cluster random and fixed effects, which, as noted earlier, tend to lower, somewhat, the absolute value of the education coefficient, while the last column reports the baseline income elasticities, evaluated at the sample mean probability.[37]

[37] For the ln variables (weight, height and sleeping rooms), the implied income elasticity is β/R_E, where β is the coefficient. For the logit dichotomous variables, the elasticity of the probability with respect to real income is $\beta(1\text{-}P)/R_E$, where P is the mean sample value (Table I).

Table VII: Product Level Estimates of the Response to Educational Attainment

	(1) baseline	(2) cluster random effects	(3) cluster fixed effects	(4) baseline Y elasticity
Radio	.153 (.001)	.149 (.001)	.134 (.001)	0.56
Television	.236 (.001)	.220 (.001)	.192 (.001)	1.21
Refrigerator	.253 (.001)	.236 (.001)	.202 (.001)	1.64
Bicycle	.056 (.001)	.077 (.001)	.078 (.001)	0.34
Motorcycle	.190 (.001)	.200 (.001)	.193 (.001)	1.47
Car	.250 (.001)	.234 (.001)	.191 (.001)	2.01
Telephone	.248 (.001)	.227 (.001)	.192 (.001)	1.77
Electricity	.228 (.001)	.235 (.002)	.216 (.001)	0.93
Tap Drinking Water	.076 (.001)	.057 (.001)	.046 (.001)	0.36
Flush Toilet	.234 (.001)	.224 (.002)	.196 (.001)	1.37
Constructed Floor	.210 (.001)	.207 (.001)	.185 (.001)	0.73
ln(Rooms/Capita)	.020 (.000)	.015 (.000)	.012 (.000)	0.17
ln Weight	.007 (.000)	.006 (.000)	.005 (.000)	0.06
ln Height	.002 (.000)	.002 (.000)	.001 (.000)	0.02
No Diarrhea	.033 (.001)	.032 (.001)	.021 (.001)	0.06
No Fever	.019 (.001)	.019 (.001)	.014 (.001)	0.05
No Cough	.006 (.001)	.008 (.001)	.005 (.001)	0.02
Alive	.059 (.002)	.059 (.002)	.046 (.002)	0.04
At School (6-14)	.200 (.001)	.171 (.001)	.151 (.001)	0.50
At School (15-24)	.148 (.001)	.135 (.001)	.111 (.001)	0.89
Working (15-24)	-.032 (.002)	-.037 (.002)	-.042 (.003)	-0.16
Working (25-49)	.020 (.001)	.028 (.001)	.025 (.001)	0.08
Birth (15-24)	-.012 (.001)	-.012 (.001)	-.007 (.002)	-0.07
Birth (25-49)	-.026 (.001)	-.024 (.001)	-.011 (.001)	-0.19
Marriage (15-24)	-.058 (.001)	-.035 (.001)	-.002 (.001)	-0.28
Marriage (25-49)	-.077 (.001)	-.064 (.001)	-.025 (.002)	-0.04

Note: The reported number is the coefficient (standard error) on household mean adult educational attainment in years, with each equation including a complete set of country x survey x region (urban/rural) dummies and the following controls: (1) consumer durables & housing: ln number of persons in the household; (2) children's health: sex, ln(1+age in months) and ln(1+age in months) squared (for all but height and weight, which are quite linear in ln(1+age)); (3) household economics: age and age squared, as well as sex for education attendance variables (all others refer to women alone). Each equation is estimated separately.

For our purposes, the main relevance of Table VII is that it establishes that each of the real consumption variables used in this paper is very significantly and, generally, quite substantially related to real income, as measured by years of education. Across the different specifications, only one coefficient (marriage of young women with cluster fixed effects) is even close to being insignificant at the 1% level. The baseline income elasticities, coupled with the standard deviation (s.d.) of mean household adult education (4.5 years) and implied s.d. of predicted incomes (4.5*.116 ≈ .5), produce substantial variation in predicted outcomes. Thus, a one s.d. movement in educational attainment produces a ln 28 percent higher relative probability of owning a radio (mean value of .574 - see Table I) and a ln 69 percent higher probability of having a flush toilet (.322). Given the early age of the subjects (0-35 months), children's weight and height move relatively less, an average of 3 and 1 percent, respectively, with a s.d. movement in educational attainment, but are, nevertheless, very significantly correlated with household incomes. The cumulative probability of survival for the average 0 to 35 month year old (mean value of .930) rises 2 percent with a s.d. movement in predicted incomes, a small apparent movement, but actually an implied fall in average cumulative mortality from .07 to .05. The probability children and youths are in school rises 25 percent (mean value of .712) and 45 percent (.304) with a s.d. movement in incomes, while the probability a young woman is working (.416) or ever-married (.431) falls by 8 percent and 14 percent, respectively. The income elasticities implied by the coefficients in the other columns can be arrived at by multiplying column (4) by the ratio of each column's coefficient to that listed in column (1).

(c) The Growth and Standard Deviation of Real Consumption

Table VIII below presents second step estimates of the growth and standard deviation of living standards using the baseline product level estimates of income elasticities and product x region x time constant terms, as described in section V above, to produce the dependent variable. In the top panel of the table, successive random effects are added, controlling for country level correlations at the product and product group level and for growth rate correlations within

Table VIII: DHS Estimates of the Growth and Standard Deviation of Living Standards
Dependent variable = urban/rural weighted country means

All products combined: $y_{pct} = a_p + g_{\sim A}*t + g_A*t + u_c + v_p*t + v_c*t + u_{pc} + u_{Gc} + e_{pct}$

$g_{\sim A}$	.035 (.006)	.037 (.002)	.038 (.006)	.038 (.006)
g_A	.032 (.005)	.035 (.002)	.033 (.005)	.033 (.005)
$\sigma[u_c]$	.743 (.073)	.708 (.072)	.714 (.072)	.707 (.074)
$\sigma[u_{pc}]$		.868 (.020)	.872 (.020)	.835 (.020)
$\sigma[v_p]$			.019 (.003)	.019 (.003)
$\sigma[v_c]$			.015 (.002)	.015 (.002)
$\sigma[u_{Gc}]$				.280 (.042)
$\sigma[e_{ct}]$	.888 (.014)	.268 (.006)	.241 (.006)	.241 (.006)

By product group: $y_{pct} = a_p + u_c + g_{\sim A}*t + g_A*t + u_c + v_p*t + v_c*t + u_{pc} + e_{pct}$

	Consumer Durables	Housing	Children's Health	Family Economics
$g_{\sim A}$	.046 (.010)	.038 (.011)	.033 (.006)	.030 (.006)
g_A	.055 (.010)	.018 (.011)	.034 (.006)	.025 (.006)
$\sigma[u_c]$	.743 (.090)	1.08 (.123)	.578 (.068)	.594 (.071)
$\sigma[u_{pc}]$	.969 (.042)	1.01 (.053)	.506 (.030)	.763 (.035)
$\sigma[v_p]$	.024 (.007)	.017 (.006)	.006 (.005)	.010 (.005)
$\sigma[v_c]$	.016 (.004)	.027 (.005)	.013 (.005)	.014 (.003)
$\sigma[e_{ct}]$	.221 (.009)	.252 (.014)	.274 (.018)	.206 (.010)

countries and products. Cumulatively, these adjustments increase the estimated growth rates, while lowering the estimated cross sectional variation in living standards. They are clearly, however, not crucial. Overall, the DHS data suggest a level of cross sectional variation consistent with that present in PWT measures of consumption per capita or per equivalent adult (between .65 and .68, in Table III earlier), but the estimated DHS non-African growth rate is close to double that recorded in PWT or UN sources, while the sub-Saharan growth rate is about 3.5 times as large. The bottom panel of the table calculates the same measures at the product group level, showing that this basic pattern is present in virtually all product group categories. Sub-Saharan growth is somewhat lower in housing (although still double the PWT growth rate) and the cross sectional variation is highest in that product group, while consumer durables show the highest growth rates. Overall, however, the same pattern is reproduced in all areas as all product groups suggest a level of cross sectional variation broadly consistent with that in PWT, but real consumption growth that is, at a minimum, double that present in the PWT and UN datasets. Sub-Saharan Africa is clearly not stagnating, but rather growing at a rate close to that of the non-African economies.

Table IX explores the sensitivity of the results to various econometric techniques and functional form assumptions. To begin, in the first column I present results where the covariance matrix of the estimated dependent variables is not incorporated in the second step GLS likelihood. As shown, this dramatically raises both the growth and standard deviation as the procedure no longer corrects for the fact that much of the cross sectional and time series variation in the dependent variable comes from the error in the first step estimates, particularly in the estimate of the income elasticity, which produces correlated expansions and contractions of the dependent variables. Turning to the second and third columns, these incorporate cluster level random and fixed effects in the first step equations used to produce the dependent variable. With somewhat smaller estimated income elasticities, on average, they expand the overall variation in the sample, producing somewhat higher estimates of the growth rate and standard deviation than the comparable numbers in the upper-right hand panel of Table VIII.[38] Continuing, the remaining four columns of the table use different functional forms (as described earlier in section

[38]Without the covariance adjustment, the random effects specification indicates non-African and African growth rates of .055 (.019) and .078 (.018), respectively, while the fixed effects specification produces the numbers .086 (.033) and .132 (.032). This illustrates the importance of incorporating the first step estimate of the covariance of the dependent variables into the MLE likelihood. I should also note that because of the lower t-statistics of the fixed effects specification, this is one of the few cases where the use of the Monte Carlo techniques to estimate the mean and covariance of the ratios of the first step variables (as discussed in the preceding section) has any effect on the estimates. Using the delta

Table IX: Sensitivity Tests $y_{pct} = a_p + u_c + g_{\sim A}*t + g_A*t + u_c + v_p*t + v_c*t + u_{pc} + u_{Gc} + e_{pct}$							
	first step logit for dichotomous variables			aternative first step function forms			
	2nd step w/out 1st step covariance	1st step cluster random effects	1st step cluster fixed effects	Probit	Weibull	Cauchy	Linear
$g_{\sim A}$	.047 (.019)	.047 (.006)	.044 (.008)	.037 (.005)	.039 (.005)	.046 (.007)	.037 (.005)
g_A	.072 (.018)	.042 (.006)	.036 (.007)	.032 (.005)	.029 (.005)	.041 (.007)	.029 (.005)
$\sigma[u_c]$	.926 (.120)	.886 (.095)	.752 (.086)	.675 (.071)	.675 (.071)	.965 (.106)	.655 (.069)
Note: each specification includes the full set of error terms (v_p, v_c, u_{pc}, u_{Gc}, e_{pct}) as in the upper right panel of Table VIII, but only g_i & $\sigma[u_c]$ are reported.							

IV) in the calculation of the first step estimates. The results are remarkably similar, with the exception of the cauchy which increases the relative change associated with differences and movements in the tails of the distribution, producing about a 1/3 increase in the estimated growth rate and standard deviation of consumption.[39]

In sum, using an estimated educational income profile (R_E) of .116, I find that the discrepancy between the DHS and the PWT appears to be concentrated in growth, with the DHS indicating 3.8 and 3.3 percent average growth of real consumption in the non-African and sub-Saharan economies, respectively, as compared to the 1.7 to 2.2 and 0.9 to 1.0 percentage growth for these country groups indicated in PWT and UN sources.[40] The cross sectional standard

method to estimate the covariance matrix, the coefficients in the third column of the table are .047 (.008), .038 (.007), and .827 (.091).

[39]Relative to the logit, probit and weibull, the cauchy has dramatically thicker tails. Hence, any given difference or movement in the mean value of a random variable in the tails (e.g. a change in the mean ownership of cars from .05 to .1) is associated with a much greater movement in the index (relative to movements around a mean value of .5).

[40]I should re-emphasize that my calculation of PWT & UN growth rates, earlier in Table III, are based upon the same country x year sample as the DHS. These figures do not represent the growth of all African and non-African economies in those data sets.

deviation of real consumption in PWT and the DHS, at about .7, is comparable. These patterns are, broadly speaking, present across all product sub-groups within the DHS data and are robust to the use of alternative functional forms or econometric specifications. A lower educational income profile would lower my estimates of growth in the DHS, but would introduce an alternative paradox, as the DHS would now indicate much less cross-sectional variation in consumption expenditure than the PWT. More importantly, adjustments of R_E will not change the fact that, in stark contrast to typical cross-national data sources, in the DHS sub-Saharan African growth in recent decades appears to be on par with that experienced by the non-African economies.

Clearly, the only way to reconcile the DHS movements in African consumption levels, which are similar to those experienced by the non-African economies (as shown earlier in Figure I), with the comparatively minimal African consumption growth reported in cross-national data sources is to introduce some form of heterogeneity in demand, allowing for the possibility that the income elasticities of demand, for the products in the DHS, are much stronger in Africa than elsewhere. To this end, in the next section I extend my methodology to allow for local variation in patterns of demand and the return to human capital. As will be seen, with this extension, I am no longer able to compare country consumption levels or calculate, meaningfully, the standard deviation of international living standards. However, the growth measures become, in a sense, more precise, as they are now calculated with local demand patterns. Thus, the "heterogeneous demand" approach provides an important robustness check on the principal, growth related, results of this paper.

VII: Methods III: Heterogeneous Demand

There is likely to be substantial variation across countries in the relationship between total consumption expenditure, product level consumption, and measurable outcomes. Two obvious examples are the demand for tap water (depending upon its cleanliness) and the relationship between nutrition and childhood weight and height (depending upon genetics). The random effects of earlier sections try to control for these, by estimating the degree to which consumption levels are correlated within product and product groups in a country, but not across all products in that country. These change the relative weight placed on repeated observations of a product across surveys in estimating levels and growth rates. These adjustments are adequate if heterogeneity manifests itself as a permanent level change in demand patterns, as motivated by the relative price effects of equation (6)' earlier. In this section I expand the treatment of heterogeneity to explicitly allow for different demographic and, especially, income effects on demand. This is purchased at some cost. The simplifying assumption of common international quasi-income elasticities of demand, used to evaluate consumption levels in previous sections, allowed the estimation of both growth rates and relative levels of living standards, much as the use of a common fixed set of international prices to weight local expenditure allows the estimation of the same in data sets such as PWT. Similarly, allowing income elasticities of demand to vary country by country is akin to weighting local expenditure using fixed country price weights, like the constant price national accounts, allowing for the international comparison of the growth rates, but not the levels, of living standards. As the principal conflict of this paper with standard sources centres on growth rates, and not relative levels, this restriction is not too onerous and allows me to establish that the most important results persist when due cognizance is made of local demand patterns.

Consider estimating the demand equation (18) earlier country by country:

$$(24)\ \ln\left(Q_{hprt}\right) = a_{prt} + b_p^c R_E^c E_{hrt} + \vec{c}_p^{\,c/} X_{hrt} + e_{hprt}$$

where the superscript c highlights the fact that the quasi-income elasticities and demographic controls are estimated country by country, while the return to education is a country-specific "known" measure. As before, the regional (urban/rural) x time (survey) product dummies can be divided by the estimated income elasticity and added to the real consumption attributable to mean education levels to produce product level country-specific measures of real consumption expenditure, which are then weighted by urban/rural population shares to produce country-level measures:

$$(25)\ \ln\left(\hat{C}_{prt}^{R^c}\right) = \frac{\hat{a}_{prt}}{\hat{b}_p^c} + R_E^c \hat{\bar{E}}_{rt} \qquad \ln\left(\hat{C}_{pct}^{R^c}\right) = \sum_{i \in c(r)} S_i \ln\left(\hat{C}_{pit}^{R^c}\right)$$

These country-specific measures, for an international panel, can then be projected in the random effects regression:

$$(26)\quad \ln\left(\hat{C}_{pct}^{R^c}\right) = a_{pc} + g*t + v_c*t + v_p*t + e_{pct}$$

The product dummies (a_{pc}) are calculated at the product x country level, explicitly recognizing that local factors influence the level of product demand. Contrasted with the comparable equation, (22), in section V earlier, one sees that these terms substitute for the global product dummies (a_p), product x country random effects (u_{pc}) and country random effects (u_c). While consumption levels are not comparable, the growth rate of country specific consumption is translated, via the local income elasticity, into comparable income equivalents, so a common international growth rate g can be calculated. Country and product level variation in growth rates is accounted for by the random effects, v_c and v_p. The stronger these are, the less relative weight is placed on countries or products with a larger number of intertemporal observations in calculating the average growth rate *g*.

As before, it should be apparent that the entire estimation procedure is linear in R_E, provided the same measure is used for all countries in equations (24) and (26). As I do not have income data for each and every one of the individual countries in my sample, I estimate (in the next section) separate African and non-African returns to education, using the data for the available countries. I then run equation (24) for each country, applying either the African or non-African R_E, and (26) for the African and non-African countries in groups. Consequently, as before, in considering my estimates of the growth rates of the African and non-African countries, the reader is free to simply modify my estimated growth rates by the ratio of his/her preferred estimate of the return to education to the African and non-African numbers that I take as "known".

VIII. Results III: Growth using Local Income Elasticities

As a preliminary, Table X below runs separate Mincerian regressions for the African and non-African countries of ln earnings from working for others on education and demographic characteristics following the specifications described in Table VI and Section VI earlier. As can be seen, the return to education appears to be higher in Africa in all formulations. As before, I instrument with the educational attainment of other household members to control for measurement error, which becomes an increasingly serious concern as additional local fixed effects are added. Comparing the last two columns of the table, the proportional attenuation bias from measurement error appears to be roughly the same for the two groups of countries, with an implied measurement standard error of 1.5 in both cases. I take the IV specification, with an

Table X: Ln Wage Regressions by Region						
		(1) survey dummies	(2) survey & rural/urban dummies	(3) cluster random effects	(4) cluster fixed effects	(5) cluster fixed effects (IV)
Africa	educ	14.00 (.281)	12.89 (.276)	12.35 (.237)	11.28 (.295)	13.86 (.944)
	age	6.38 (1.24)	6.40 (1.20)	6.37 (1.07)	5.32 (1.26)	5.13 (1.54)
	age2	-0.06 (.016)	-0.06 (.016)	-0.06 (.014)	-0.05 (.017)	-0.04 (.020)
	sex	-4.27 (3.75)	-5.58 (3.66)	-6.27 (2.58)	-6.10 (2.99)	-3.02 (3.81)
	N	8041	8041	8041	8041	5897
~Africa	educ	10.27 (.176)	9.80 (.169)	9.52 (.148)	8.67 (.179)	10.33 (.534)
	age	4.23 (.795)	4.17 (.791)	4.62 (.746)	5.11 (.830)	4.97 (.953)
	age2	-0.04 (.011)	-0.04 (.011)	-0.05 (.010)	-0.05 (.011)	-0.05 (.013)
	sex	-54.8 (1.91)	-55.4 (1.91)	-55.3 (1.86)	-53.9 (2.04)	-55.5 (2.31)
	N	14955	14955	14955	14955	12521

For notes and details on variable construction see Table VI and Appendix I.

Table XI: Cross-Country Heterogeneity of Logit or Regression Coefficients on Household Educational Attainment							
	Mean Country Coefficient	Standard Dev. of Coef.	N		Mean Country Coefficient	Standard Dev. of Coef.	N
Radio	.162 (.006)	.043 (.004)	55	Electricity	.235 (.012)	.084 (.009)	53
Television	.252 (.009)	.063 (.006)	55	Tap Water	.091 (.009)	.066 (.007)	55
Refrigerator	.264 (.009)	.067 (.007)	54	Flush Toilet	.248 (.008)	.058 (.006)	53
Bicycle	.059 (.010)	.071 (.007)	55	Cons. Floor	.205 (.010)	.071 (.007)	54
Motorcycle	.161 (.012)	.086 (.009)	55	ln(Rms/Capita)	.016 (.002)	.012 (.001)	50
Car	.244 (.008)	.057 (.006)	53				
Telephone	.270 (.010)	.072 (.008)	52	At School (6-14)	.208 (.009)	.066 (.007)	56
				At School (15-24)	.163 (.009)	.068 (.007)	55
ln Weight	.007 (.000)	.002 (.000)	51	Working (15-24)	-.009 (.007)	.044 (.005)	49
ln Height	.002 (.000)	.001 (.000)	51	Working (25-49)	.052 (.010)	.067 (.007)	49
No Diarrhea	.035 (.004)	.023 (.003)	55	Birth (15-24)	-.014 (.003)	.014 (.003)	56
No Fever	.020 (.003)	.019 (.003)	55	Birth (25-49)	-.033 (.004)	.023 (.003)	56
No Cough	.005 (.003)	.023 (.003)	55	Marriage (15-24)	-.050 (.007)	.050 (.005)	56
Alive	.057 (.005)	.030 (.004)	56	Marriage (25-49)	-.089 (.008)	.058 (.006)	56

Notes: N = # of country-level estimating equations. Numbers in parentheses are standard errors. Mean and standard deviation estimated taking into account the 1st step standard errors of the coefficients on household educational attainment, as described in the text.

estimated return to education of .1386 in Africa and .1033 outside of Africa, as the basis for the analysis further below.[41]

Table XI describes the strong heterogeneity across countries in demand patterns. Taking the estimated demand response of each product to mean household educational attainment in each country as data, I regress these dependent variables on a constant, incorporating the first step covariance matrix in the likelihood, as described earlier in section III. The coefficients reported in the table are that of the constant (mean country coefficient) and the standard error of the regression (standard deviation of the coefficient).[42] As can be seen, the standard deviations are very large relative to the mean values of the coefficients, reflecting the degree of heterogeneity. To cite just one example, while the demand for tap water is strongly positively associated with educational attainment in the world as a whole (mean coefficient = .091), it is very strongly negatively associated with educational attainment in the Dominican Republic (coef. = -.10), where tap water is known to be frequently contaminated.

Turning to growth measures, Table XII presents separate estimates of growth in the African and non-African economies based on equation (26) in Section VII. It is immediately apparent that the considerable heterogeneity in demand patterns described above has little effect on the results. Focusing on the baseline logit formulation, African growth is now seen to be

---

[41]As shown in the table, women appear to face a negligible discount in the labour market in sub-Saharan Africa. This is a place where selectivity bias is likely to play a major role and, indeed, adjustments along this dimension yield the expected results. When I estimate the wage equation formulation of column (2) jointly with a labour participation equation using marital and pregnancy status as independent determinants of participation (as described in the footnote associated with selectivity bias in Table VI), the woman's discount rises to 29% in Africa, while remaining at 59% for the non-African economies. However, as before, the educational income profile, at 13.5 and 9.8 within and outside Africa, respectively, is largely unchanged.

[42]These are not, exactly, the means and standard deviations of the dependent variables, as the constant term is adjusted for weighting based upon the precision of each estimate and the standard error of the regression is reduced by the MLE's recognition that part of the variation in the dependent variables is simple estimation error.

Table XII: Growth Measures Based on Local Demand Patterns $y_{pct} = a_{pc} + g_i*t + v_p*t + v_c*t + e_{pct}$				
	Baseline Logit	1st step cluster random effects	1st step cluster fixed effects	2nd step w/out 1st step covariance
$g_{\sim A}$ g_A	.034 (.005) .038 (.005)	.042 (.006) .042 (.007)	.053 (.008) .048 (.009)	.047 (.034) .166 (.060)
	Probit	Weibull	Cauchy	Linear
$g_{\sim A}$ g_A	.033 (.005) .036 (.005)	.032 (.005) .034 (.004)	.040 (.006) .044 (.007)	.036 (.006) .031 (.004)
Note: African and non-African growth rates estimated in separate equations, each containing a full set of product x country dummies and random effects for regional product and country level variation in growth rates (v_p & v_c).				

somewhat higher than previously estimated in Table VIII (.038 vs .033) and non-African growth somewhat lower (.034 vs. .038). This, however, is entirely due to the use of a higher R_E for Africa and a lower R_E for the non-African economies in this section. With a common R_E of .116, as used previously, the growth rates are virtually identical to those estimated earlier. As before, estimates with random and fixed effects yield somewhat higher growth rates, estimates without adjustment for the 1st step covariance are extraordinary and nonsensical (methodologically and practically), and alternative functional forms yield similar numbers with, once again, the cauchy producing the highest estimates of growth.[43] There is, without a doubt, considerable heterogeneity across countries in demand patterns, just as there is considerable heterogeneity in

[43]As noted earlier in section III, this one of the few cases where use of Monte Carlo techniques to estimate the mean and variance of the dependent variable (calculated as a ratio of multivariate normals) yields different results than those arrived at by using the (inaccurate) delta method. With delta method means and covariances, the growth rates for 1st step cluster fixed effects are .048 (.007) and .046 (.007) for the non-African and African economies, respectively, while the growth rates for the baseline without 1st step covariance are .082 (.119) and .018 (.053). The latter case is particularly interesting, because the baseline calculations with the covariance matrix included (the upper left-hand panel of the table) are identical with the Monte Carlo and delta method covariance matrices. Without the covariance matrix, the wild mean values of equations estimated, in the restricted sample country-level equations, with little accuracy dominate the estimates. With the covariance matrix implicit weighting, of either method, these play no role.

prices, but this averages out completely, and cannot eliminate the surprisingly high growth, particularly for sub-Saharan Africa, indicated by the DHS data.

IX. Inequality: The Ln of the Mean and the Mean of the Ln

A conceptual difference in measurement, associated with inequality, explains part of the difference between my DHS based estimates and those of the PWT. Conventional measures of real living standards calculate the average consumption per capita, i.e. total aggregate real consumption divided by the number of persons or equivalent adults. When the growth and cross-national dispersion of these measures is examined, the ln is usually taken, so they may reasonably be termed the "ln of the mean". In my use of the DHS data to produce estimates of real living standards, I calculate the ln income equivalent of constant terms estimated in household demand functions which are linear in ln income (proxied by years of education). These ln income equivalent constant terms are then aggregated using regional weights. Thus, my procedure can reasonably be described as estimating the "mean of the ln".

From Jensen's inequality, we all know that the mean of the ln is less than the ln of the mean. What is more relevant, however, is that the gap is related to the degree of dispersion of the random variable; in our case, the dispersion of real consumption. To make things concrete consider, the case where ln household real expenditure per adult, $\ln(C^R)$, is ln normally distributed with mean μ and standard deviation σ. In this case we have the well known result:

$$(27) \quad E[\ln(C^R)] = \mu \qquad \ln[E(C^R)] = \mu + .5\sigma^2$$

While (27) is exactly correct if real expenditure is ln normally distributed, it is also true as a (surprisingly accurate) second order approximation for any distribution of living standards.[44] More generally, as many readers are no doubt aware, the difference between the ln of the mean and the mean of the log is none other than Theil's mean log deviation index of inequality. Thus, it is fair to say that conventional measures of living standards, such as the PWT, are actually a mixture of the average of ln real living standards and the dispersion, or inequality, of the same.

It is quite easy to develop ln of the mean equivalent DHS consumption measures using the estimates presented in previous sections. For any product level measure of ln real regional consumption, estimated using common international quasi-income elasticities (equation (20) above) or using heterogeneous country specific demand patterns (equation (25) above), mean of the ln and ln of the mean country measures can be calculated by either weighting the ln measure or the exponential of the ln measure:

$$(28)\quad E[\ln(C^R_{pct})] = \sum_{i \in r(c)} S_i \ln(\hat{C}^R_{pit}) \qquad \ln[E(\ln(C^R_{pct}))] = \ln\left(\sum_{i \in r(c)} S_i e^{\ln(\hat{C}^R_{pit})}\right)$$

These measures can then be projected in a random effects regression to produce estimates of the growth and cross-national standard deviation of real consumption, as described in previous sections.[45]

---

[44]Thus, take the variable x distributed with mean μ and variance σ^2. Using a second order approximation:

$$E(e^x) \cong e^{\mu} + .5e^{\mu}\sigma^2 \quad \text{so that} \quad \ln[E(e^x)] \cong \mu + \ln(1+.5\sigma^2) \cong \mu + .5\sigma^2$$

where the last approximation follows if σ^2 is small. The distribution of household educational attainment in the DHS data is horribly left skewed and bounded from below by 0, i.e. by no means normal. Nevertheless, when I calculate ln mean household income, i.e. ln[E(exp(.116*E_i))], using the approximation listed above and using the actual distribution of educational attainment, I get virtually identical numbers.

[45]As specified, (28) assumes there is no variation in ln expenditure within urban/rural regions. In a separate paper I expand the methodology to allow the calculation of within region inequality (across clusters and within clusters), allowing for a fuller description of overall levels of inequality. For the

Table XIII: DHS Estimates of the Mean of the Ln and the Ln of the Mean					
		Country Level		Urban	Rural
		ln[E(C)]	E[ln(C)]	E[ln(C)]	E[ln(C)]
Homogeneous Demand	$g_{\sim A}$	.034 (.005)	.038 (.006)	.023 (.005)	.042 (.006)
	g_A	.032 (.005)	.033 (.005)	.026 (.005)	.033 (.006)
	$\sigma[u_c]$	.671 (.070)	.707 (.074)	.495 (.055)	.640 (.069)
Heterogeneous Demand	$g_{\sim A}$	.031 (.005)	.034 (.005)	.022 (.004)	.035 (.006)
	g_A	.035 (.005)	.038 (.005)	.028 (.005)	.037 (.005)

Notes: Homogeneous demand follow the specification of the upper right hand panel of Table VIII: $y_{pct} = a_p + g_{\sim A}*t + g_A*t + u_c + v_p*t + v_c*t + u_{pc} + u_{Gc} + e_{pct}$. Heterogeneous demand follow the specification of Table XII: $y_{pct} = a_{pc} + g_i*t + v_p*t + v_c*t + e_{pct}$. In the case of heterogenous demand, the equations are estimated separately for the African and non-African economies.

Table XIII above compares estimates of the growth and standard deviation of living standards based upon the ln of the mean and the mean of the ln. In the upper panel, I use the specification of section VI earlier, with common international demand coefficients, to estimate both the growth rate and the cross-national standard deviation of consumption expenditure. The first column presents ln of the mean estimates, the PWT equivalent measure, while the second column reproduces the mean of ln estimates of Table VIII earlier. As can be seen, moving to the standard measurement concept lowers the estimated growth rates, particularly for the non-African economies. The principal force behind this difference is presented in the third and fourth columns of the table, which calculate separate mean of the ln growth rates for urban and rural areas.[46] As can be seen, rural growth is much more rapid. Given the relative poverty of rural

purposes of moving the DHS growth rates towards PWT, however, the main effect revolves around the simple gross differences between the urban/rural means highlighted above.

[46]In this case the dependent variable is simply the urban or rural estimates of product level real consumption $\ln(Y^R_{pit})$ without any need for urban/rural weighting.

areas,[47] this implies a reduction in inequality, and, hence, slower ln of the mean measures of growth.[48] The lower panel of the table uses the heterogeneous (local coefficients) demand model of section VIII and finds, again, that the movement to the PWT concept lowers growth rates. With local demand coefficients, the non-African growth rate of 3.1 percent is within striking distance of the 1.7 - 2.2 percent per annum suggested by international sources, but African growth, despite the downward adjustment to 3.5 percent, remains an eye-opening three and a half to four times the .9 to 1.0 percent indicated by PWT and UN sources.

X. Conclusion

Figures II and III summarize the points of agreement and disagreement between the DHS and the PWT. Figure II begins by graphing the estimates of relative country consumption expenditure levels in the year 2000 suggested by the two datasets.[49] Clearly, although the two estimates often differ substantially, they are highly correlated. In both datasets, the African economies make up most of the countries in the poorer lower-left hand quadrant of the figure and very few of the countries in the rich upper-right hand quadrant. A dummy variable for Africa's relative poverty has a value (s.e.) of -.677 (.150) in the PWT and a somewhat worse, but not

[47]I can estimate the average urban-rural gap by running the random effects equation of the table using the differenced urban-rural product level consumption measures, (20) earlier, as the dependent variable and treating the product constant (a_p) as a random effect, to allow the estimation of a common constant. I find the average urban-rural gap to be .797 (.140) outside of Africa and 1.028 (.139) in sub-Saharan Africa.

[48]The adjustment for inequality also lowers the cross-national standard deviation of incomes, reflecting the fact that inequality is found to be greater in poorer countries. As noted in the footnote above, I explore these issues more fully in a separate paper.

[49]For the DHS, I run the specification of the ln of the mean consumption in the upper-left hand panel of Table XIII, but with a fixed effects calculation of country dummies u_c. Similarly, the PWT use a fixed effects version of the equation in Table III for ln consumption per equivalent adult. The data plotted in the figure are the demeaned fixed effects.

Figure II: Relative Real Consumption

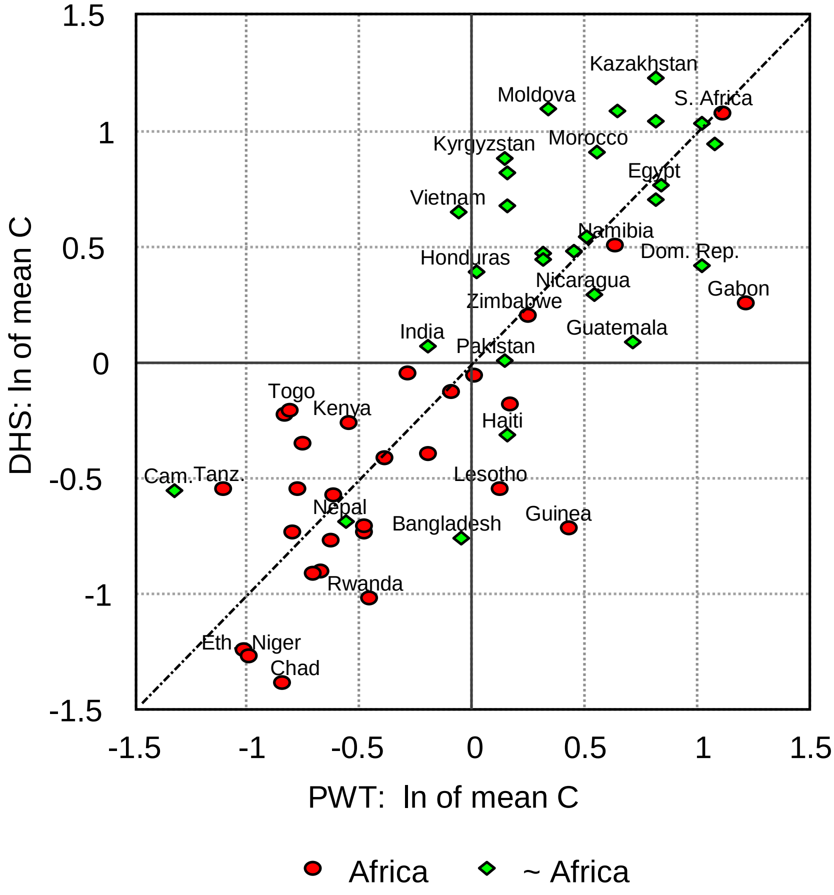

significantly different value, of -.901 (.145) in the DHS.[50] The DHS point estimate (s.e.) of the overall standard deviation of ln mean country consumption, at .671 (.070), is virtually identical to the PWT figure of .651 (.062).[51] In regards to levels of consumption, the DHS and PWT are in broad agreement.

[50]Estimated by augmenting the country random effects regressions of the second column of Table III and the upper left-hand panel of Table XIII with a dummy variable for the relative level of African consumption in the base year.

[51]See Tables XIII and III earlier. The random effects calculation of the standard deviation of u_c in those tables, the fixed effects specifications used to estimate the dummies for Figure II, and the random effects with African dummies used to derive the point estimates reported above are all, in essence, similar calculations. As noted in section III, if one wants to proceed sequentially, by first estimating country fixed effects and then using these to calculate a standard deviation or the relative poverty of Africa, one

Figure III highlights the substantial inconsistency, in growth rates, between the data of the PWT and the DHS. The various panels project the country demeaned dummies (level estimates) for each country x survey time period against the country demeaned year. This is the variation that identifies the growth rates estimated with each data source.[52] In the upper panel, we see that although the PWT suggest slower growth in non sub-Saharan African countries than is indicated by the DHS, there is enough dispersion around the mean growth, i.e. uncertainty about the estimates, to make the reported difference unimpressive.[53] Turning to the lower panel, however, we see that sub-Saharan growth in PWT is negligible, while in the DHS data it is strong, clearly significant and on par with that present in non-African countries. There is simply too much of an upward trend in the measured consumption of the DHS sub-Saharan countries to be consistent with the utter stagnation implied by the PWT, and other cross-national, data for the region. African consumption is growing faster than cross-national data sources, drawing on a mixture of country national accounts reports and ad hoc extrapolations and interpolations, indicate.

---

should, to be formally correct, incorporate each step's covariance matrix into the GLS of the next step and, eventually, calculate complicated three-step standard errors. However, for the reader's information, the standard deviation of the fixed effects dummies in Figure II are .704 (DHS) and .656 (PWT) and the differences in the average of the African and non-African dummies are -.911 (DHS) and -.680 (PWT), i.e. virtually identical to the estimates reported above.

[52]For PWT, the country x survey time period dummy is simply the reported data. For the DHS, I run the specification $y_{pct} = a_p + u_{ct} + v_p*t + u_{pc} + u_{Gc} + e_{pct}$ (the regional time trend and country growth random effects are no longer identified), estimating the u_{ct} as fixed effects. In the figure, I remove the 14 countries with only one (time period) observation, as the residuals are automatically zero for both the DHS & PWT. The lines drawn in the panels are those implied by the African and non-African growth estimates of the upper left hand panel of Table XIII and the second column of Table III. Those regressions, of course, include additional trend random effects and, in the case of the DHS, the first-stage covariance matrix in the GLS, so the regression lines are not simply the demeaned y variable projected against the demeaned x variable; but, as the reader can see, they are reasonably close to what that would be.

[53]To be consistent with Figure II, I use the homogenous (global) demand equations of section V and VI. As noted earlier, the heterogenous (local) demand model of sections VII and VIII produces lower DHS estimates of growth for the non-African countries, bringing them more in line with PWT and the UN.

Figure III: Deviations from Country Means

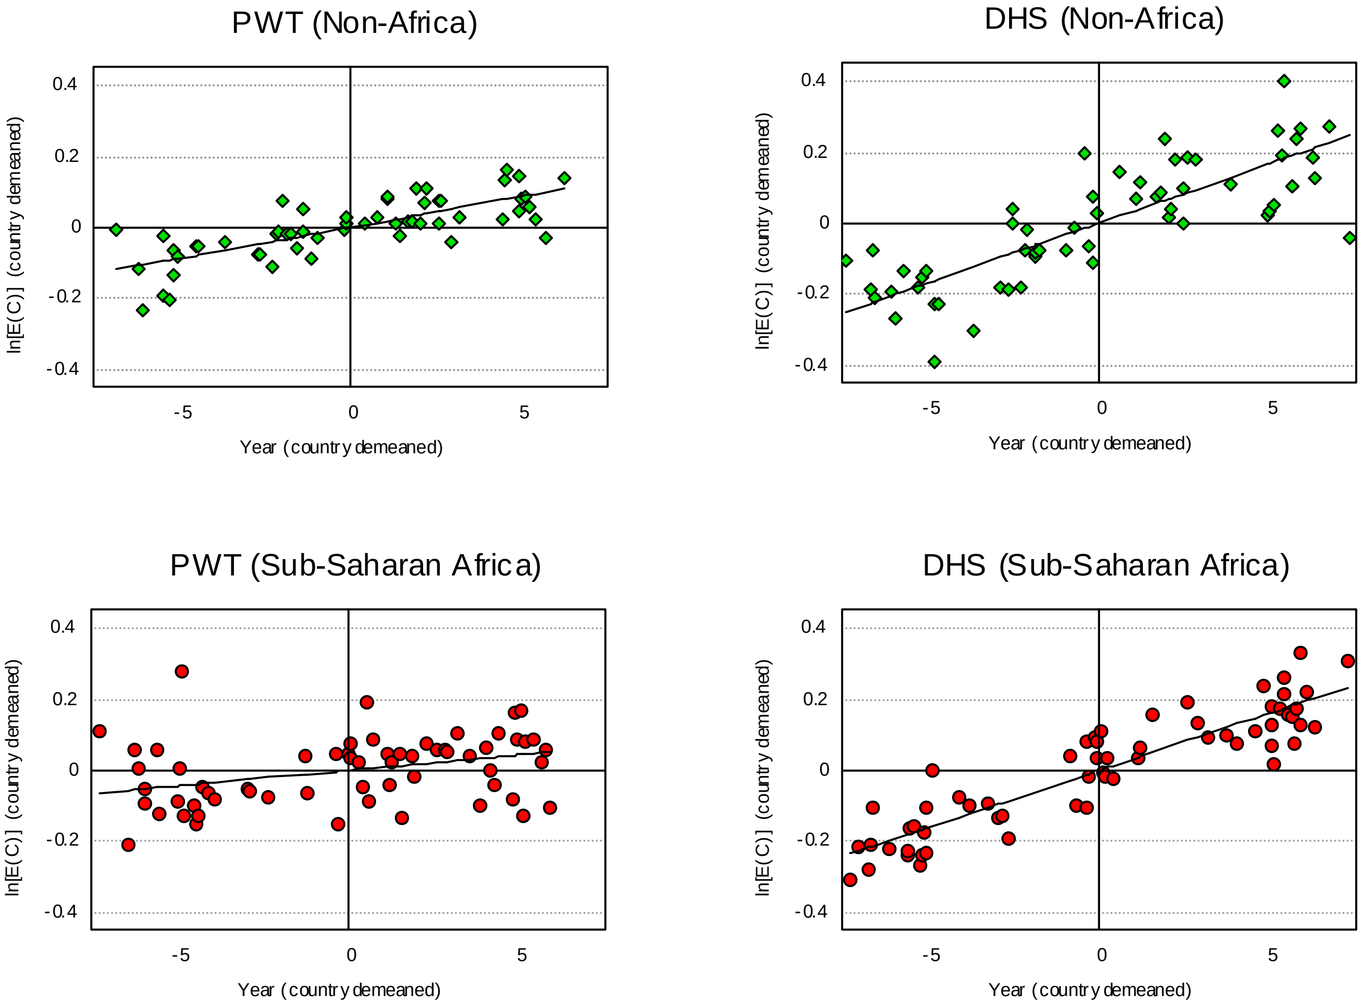

Demographic and Health Survey data on the consumption of consumer durables and housing, children's health and mortality, the schooling of youth and the allocation of women's time between marriage & childbirth and market activity, indicate that since 1990 real material consumption in sub-Saharan Africa has been rising at a rate three and half to four times that recorded by international data sources such as the PWT and UN, and on par with the growth taking place in other regions of the world. This is a miraculous achievement, given that the very real ravages of the AIDS epidemic have deprived families of prime working age adults, burdened them with medical and funeral expenses, orphaned their school age children and directly and adversely affected the health of their infants. And yet, the overall health and mortality of children is improving, their school attendance is rising, and family consumption of a variety of material goods is growing at a rapid rate. Notwithstanding these heartening trends, it is important to keep in mind that the DHS data also indicate that Africa is much poorer than other developing countries, with levels of ln consumption 90 percent lower than those enjoyed by the other developing countries in the DHS sample. For all its tragic difficulties, sub-Saharan Africa is not being left further behind by the rest of the world. It remains, nevertheless, very much behind.

XI. Bibliography

Butler, J.S. and Robert Moffitt. "A Computationally Efficient Quadrature Procedure for the One Factor Multinomial Probit Model." Econometrica 50 (1982): 761-764.

Chamberlain, Gary. "Analysis of Covariance with Qualitative Data." Review of Economic Studies 47 (1980): 225-238.

Chinapah, Vinayagum, E.L. Mostafa H'ddigui, Anil Kanjee, et al. With Africa for Africa: Towards Quality Education for All. Pretoria, South Africa: Human Sciences Research Council, 2000.

"Data Appendix for a Space-Time System of National Accounts: Penn World Table 6.1 (PWT 6.1). February 2008. At http://pwt.econ.upenn.edu/Documentation/append61.pdf.

Heston, Alan. "A brief review of some problems in using national accounts data in level of output comparisons and growth studies." Journal of Development Economics 44 (1994): 29-52.

Psacharopoulos, George. "Returns to Investment in Education: A Global Update". World Development 22 (1994): 1325-1343.

Slesnick, Daniel T. "Are Our Data Relevant to the Theory? The Case of Aggregate Consumption." Journal of Business and Economic Statistics 16 (January 1998): 52-61.

Strauss, J.P. Results of the Monitoring Learning Achievement Project. Bloemfontein, South Africa: Research Institute for Education and Planning, 1999.

Young, Alwyn. "In Sorrow to Bring Forth Children: Fertility Amidst the Plague of HIV. Journal of Economic Growth 12 (December 2007): 283-327.

XII: Appendix I: Demographic and Health Survey Data

Table A1 below lists the DHS surveys used in the paper. The DHS survey codes corresponding to the living standard variables listed in Table I above are ("hv" variables come from the household file, all others from the women's file):

> Radio (hv207), television (hv208), refrigerator (hv209), bicycle (hv210), motorcycle (hv211), car (hv212), telephone (hv221), electricity (hv206), tap drinking water (hv201), flush toilet (hv205), constructed floor (hv213), sleeping rooms (hv216), weight (hw2), height (hw3), diarrhea (h11), fever (h22), cough (h31), alive (b5), attending school (hv121 or hv110 if unavailable), working (v714), gave birth past year (v209), ever married (v502).

All "don't know" or "missing" responses are dropped from the sample. Some variables are recoded into broad dichotomous 0/1 categories, or to correct survey anomalies and differences, as follows:

> Constructed floor: hv213 <= 13 (dirt/sand/dung) = 0, otherwise (cement/wood/tiles/etc) = 1; Chad 2004 hv213 = 12 (palms/weaves) coded as 1; Bangladesh combines bamboo/earth as a single "natural floor" category, which I code as 0. Flush toilet (including septic tanks): hv205 < 21 = 1 (private/shared distinction not universal across surveys, and not used), otherwise (pit/latrine/bush/etc) = 0; Benin 1996, Togo 1998 and Indonesia 2002 removed (emphasis on private/public and covered/uncovered, no clear "flush" distinction in question); Brazil 1991 & 1996 DHS coding confused (pits coded as toilets, septic tanks as latrines) & recoded using survey documentation; Cambodia 2000 recode latrines connected to septic tanks or sewers as septic tanks. Tap drinking water: hv201 < 21 = 1 (tapped or piped, distinction about location in or out of residence not consistent across surveys and not used), otherwise (well/stream/lake/etc) = 0; Bolivia 1994 "neighbour" defined as piped or unspecified (DHS codes 13 and 14), recode unspecified as not tap or piped (my code 0). Diarrhea, fever and cough in past 2 weeks: yes answers 1 or 2 coded as 0 (DHS extra distinction between past 24 hours and past 2 weeks not universal across surveys and not used), 0 (no) coded as 1. Gave birth past year: one or more births coded as 1, none coded as 0. Marital status: currently and formerly coded as 1, never coded as 0.

Conditioning/demographic variables (see Table VII) are constructed as follows:

> Ln number of household members (number of hvidx household records); young children's sex (b4) and age in months (v008-b3); youth's sex (hv104) and age (hv105); married women's age (v012).

Because of changes in the coverage of DHS survey questionnaires over time, samples are restricted to generate consistent samples, as follows:

> Children's health variables: children aged 35 months or less (i.e. born within 35 months of the survey). Women's fertility and work variables: currently married women only.

Table AI: DHS and Associated Surveys Used in the Paper			
Benin	1996*, 2001, 2006	Bolivia	1994*, 1998*, 2003
Burkina Faso	1992, 1998, 2003	Brazil	1991, 1996
Cameroon	1991, 1998, 2004	Colombia	1990, 1995*, 2000, 2005
Cen. Af. Rep.	1994*	Dom. Rep.	1991, 1996*, 1999, 2002
Chad	1996*, 2004	Guatemala	1995*, 1998*
Comoros	1996*	Guyana	2005
Congo	2005	Haiti	1994, 2000, 2005
Cote D'Ivoire	1994, 1998, 2005	Honduras	2005
Ethiopia	2000, 2005	Nicaragua	1997*, 2001
Gabon	2000	Paraguay	1990
Ghana	1993, 1998*, 2003	Peru	1992, 1996*, 2000, 2004
Guinea	1999, 2005		
Kenya	1993, 1998, 2003	Bangladesh	1993, 1996, 1999, 2004
Lesotho	2004	Cambodia	2000, 2005
Madagascar	1992, 1997*, 2003	India	1992, 1998, 2005
Malawi	1992, 2000, 2004	Indonesia	1991, 1994, 1997, 2002
Mali	1995*, 2001, 2006	Nepal	1996*, 2001, 2006
Mozambique	1997*, 2003	Pakistan	1990
Namibia	1992, 2000	Philippines	1993, 1998*, 2003
Niger	1992, 1998, 2006	Vietnam	1997, 2002
Nigeria	1990, 1999*, 2003		
Rwanda	1992, 2000, 2005	Armenia	2000, 2005
Senegal	1992, 2005	Egypt	1992, 1995*, 2000, 2003, 2005
South Africa	1998*	Kazakhstan	1995, 1999
Tanzania	1992, 1996, 1999, 2003, 2004	Kyrgyz Rep.	1997
Togo	1998*	Moldova	2005
Uganda	1995*, 2000, 2006	Morocco	1992, 2003
Zambia	1992, 1996*, 2001	Turkey	1993, 1998*, 2003
Zimbabwe	1994*, 1999, 2006	Uzbekistan	1996

Notes: Years denote date when survey began; data collection often continues into the following year.
(*) Surveys with wage income data

For the wage regressions in Table VI, I restrict myself to female and male individuals aged 25 and above reporting that they work for others (v719 or mv719 = 2, "m" denotes the male questionnaire). Annual earnings are constructed from v736/mv736 data, with the earnings of individuals reporting annual, monthly and weekly wages multiplied by 1, 12 and 50, respectively (individuals reporting an hourly or daily wage, numbering about 1/5 of those working for others and reporting wage data, are dropped from the sample). As I have painstakingly recoded all the educational data for the household files, but have not done the same for the male and female questionnaires, I get individual age and educational characteristics by merging the individual files (which contain the earnings data) with the household files using the individual id numbers, eliminating cases where the individual's sex does not match across the two files or there is a discrepancy of more than 2 years in the reported age (roughly 7 percent of cases that meet the other wage sample eligibility criteria).

Employment, schooling and marital status pose special problems. On women's employment, variation in the question form has dramatic effects on average responses. The standard questionnaire first asks women if, apart from housework, they are currently working and then follows up with a question that explains that women may work in a variety of ways (for cash or in kind, selling things, in their businesses, on farms or in the family business) and asks the respondent if she is currently doing any of these. The combination of these two questions form the basis for DHS code v714. An occasional third question on whether the woman has done any work in the past 12 months then produces v731. The problem is that many DHS surveys vary this pattern, omitting the first or second of the two part v714 question, inserting the words "last week" into one or both of these questions, omitting the preliminary v714 questions in their entirety (but including the v731 question), and even modifying the questions to focus on working for cash only. When compared across survey years for individual countries, these changes produce very large variation in average employment rates. Consequently, I restrict my measure to v714 and only those surveys where the two-part question is asked in its standard form.

On schooling, some questionnaires ask whether the household member attended school in the past year (hv121) and others whether the household member is currently in school or still in school (hv110). The form of this question does not seem to be important, as the differences within surveys where the two questions overlap and between surveys when the questions change are small. Consequently, I take hv121 when it is available, and use hv110 as a reasonable substitute when it is not. The main problems that arise in the educational data are that (1) in some surveys individuals who, when questioned on educational attainment, say they have never been to school are automatically coded as not currently attending school, whereas in other surveys they are not; (2) the educational attendance question is generally restricted to individuals 6 to 24, but in some surveys the age range is further restricted, while those who were not asked the question are automatically coded as not attending. I solve these problems by coding all individuals whose educational attainment is listed as having never attended school as not currently attending and, in cases where problem (2) arises for 6 year olds only, coding all 6 year olds as missing. For the Indian surveys, problem (2) arises for individuals older than 14, 17 or 18 (depending on the survey). In effect, for the age group 15-24 India's education data is restricted to individuals 15 to 17 or 15 to 18, which makes India non-comparable with the other countries in my sample. Consequently, I eliminate India from the sample for this variable. In the case of the few surveys with missing data for 6 year olds, I deem that the age controls and the existence of data for the remainder of youths aged 6-14 allow me to keep them in the sample.

On marital status (never vs currently/formerly), this is reported in the women's question module which, in some surveys, is restricted to ever-married women. To code never-married women for these surveys, I begin by identifying the additional eligibility criterion for the female survey (usually "slept last night", rarely "usual resident", but the two variables are extraordinarily correlated). I then code all women in the household file meeting the additional eligibility criterion who are also listed as "not eligible" for the women's questionnaire as "never married", and merge these records with the marriage data from the women's question module. The marital status of women who do not meet the additional eligibility criterion is uncertain (they are excluded from the female survey even if they are married), so they are dropped from the marital status sample.

Finally, I turn to educational attainment. The DHS questionnaires ask respondents for their educational attainment, measured as grade level achieved, not the number of years attended.[54] The DHS "recode" takes this raw data, converts it into a broad categorical variable (hv106 = none, primary, secondary, tertiary), a measure of years at that level (hv107), and total years of attainment (hv108). Unfortunately, the procedures used by programmers to generate these conversions over the years have varied, with, for example, the number of years of education falling in each hv106 category varying even within countries. Most fundamentally, there are extraordinary errors and inconsistencies in reaching the final years of attainment (hv108), with, to cite some examples, those responding "don't know", a code of 8 in many surveys, credited 8 years of education; reaching tertiary education (not counting years there) being credited anything from 10 to 19 years base (sometimes, within the same country); upper secondary systems that require 10 formal levels to reach being coded as 6 years; etc. Working with the DHS questionnaires, original "raw" non-recode data generously provided by the DHS programmers, and summaries of educational systems and their history found on websites hosted by UNESCO, education.stateuniversity.com, jstor, and the education ministries of different countries, I have recoded all the educational attainment data to represent years of formal attainment within each country's educational ladder, taking the level of entering 6 year olds as the starting point. In cases where systems change over time (e.g. an old system primary lasted 6 years and a new system primary lasts 8 years, so "completed primary" has different meanings), I use the timing of institutional reform, an individual's birth cohort, and sample information on the distribution of years of attainment by age group (e.g. those with uncompleted primary up to a certain birth cohort indicate no more than 6 years) to impute an appropriate estimate of years of completed education to different birth cohorts.

[54]This is a more accurate measure of attainment, as grade repetition is quite common (see, for sub-Saharan Africa, Chinapah et al 2000 and Strauss 1999.