NBER WORKING PAPER SERIES

THE AFRICAN GROWTH MIRACLE

Alwyn Young

Working Paper 18490
http://www.nber.org/papers/w18490

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2012

I am grateful to Chad Jones, Pete Klenow, Ben Olken and anonymous referees for very helpful comments and to Measure DHS for making their data publicly available The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The African Growth Miracle
Alwyn Young
NBER Working Paper No. 18490
October 2012
JEL No. O47

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 3.4 to 3.7 percent per annum, i.e. three and a half to four times the rate indicated in international data sets.

Alwyn Young
Dept. of Economics
London School of Economics
Houghton St.
London WC2A 2AE
UNITED KINGDOM
and NBER
a.young@lse.ac.uk

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 provided real income estimates for 45 sub-Saharan African countries, in 24 of those countries it did not have 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 unquantifiable uncertainty. As noted by Heston (1994), consumption measures for most developing countries are derived as a residual, after subtracting the other major components of expenditure from production side estimates of

[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 were used to extrapolate the price studies of benchmark countries to non-benchmark economies. This problem has been alleviated somewhat with the 2005 ICP worldwide study of prices that informs PWT 7.0. As I show further below, the updating of PWT data in this fashion moves its level estimates systematically closer to my results.

[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).

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 log 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.

The principal result of this paper is that real household consumption in sub-Saharan Africa is growing between 3.4 and 3.7 percent per annum, i.e. three and a half to four times the 0.9 to 1.1 percent reported in international data sources. I find that the growth of consumption in non sub-Saharan economies is also higher than reported in international sources, but the difference here is much less pronounced, with growth of 3.4 to 3.8 percent, as opposed to the 2.0 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 the Penn World Tables, although there are substantial differences for individual countries.

I follow the lead of scholars such as Becker, Philipson and Soares (2005) and Jones and Klenow (2011) and take a broader view of consumption than is typically used in the national accounts, including health outcomes and the use of family time. These elements, however, do

not explain the discrepancy between my estimates and international sources. I find the real consumption equivalent of health and family time to be growing about as fast or slightly slower than the average product, so their removal leaves the main results unchanged. In general, I show that the results are not unusually sensitive to the exclusion of any particular product, while a narrow focus on the slowest growing product group of all (housing) still produces sub-Saharan growth estimates which are double those of international sources.

I begin in Section II below by describing the DHS data. Section III then presents an intuitive introduction to my method, describing how I convert data on real product consumption into money metric real consumption equivalents by dividing them by the Engel curve coefficients estimated off of household micro data. Section IV provides a more formal exposition, while Section V applies the technique to the DHS data, producing the results outlined above. The analysis of Section V imposes the simplifying assumption that a single Engel curve equation approximates global demand for a product. I relax this in Section VI, estimating Engel curves country by country, and show that the growth results are unchanged. Section VII concludes.

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 1.6 million households in 56 developing countries, as listed in Appendix A. The occasional nature of the DHS surveys means that I have an unbalanced panel with fairly erratic dates. Thus, I will not be able to meaningfully report a full set of country specific growth rates for the past two decades. I can, however, divide the sample into sub-Saharan and non sub-Saharan countries and calculate the average growth rate of each group during the period covered by the data (1990-2006). This is what I do further below.

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 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) and removed surveys with inconsistent question formats (in particular, regarding labour force participation). Appendix A lists the details.[3]

[3]The cleaned data files and all of the programmes used to produce the results of this paper are available on my website http://personal.lse.ac.uk/YoungA/. The original data are available at www.measuredhs.com.

Table I: DHS Real Living Standard Measures by Category					
	N	Mean		N	Mean
Ownership of Durables			Housing Conditions		
Radio	1549722	.573	Electricity	1526536	.530
Television	1569789	.406	Tap Drinking Water	1561296	.451
Refrigerator	1465668	.249	Flush Toilet	1441519	.323
Bicycle	1481982	.296	Constructed Floor	1392545	.599
Motorcycle	1423388	.103	log # Sleeping Rooms per Person	709399	-.927
Car	1452204	.066			
Telephone	1127789	.172			
Children's Nutrition and Health			Household Time and Family Economics		
log Weight (100g)	465085	4.44	Attending School (age 6-14)	1916473	.712
log Height (mm)	454582	6.59	Attending School (age 15-24)	1219551	.340
Diarrhea	586536	.201	Working (women age 15-24)	191822	.412
Fever	575492	.323	Working (women age 25-49)	579082	.551
Cough	582544	.342	Gave Birth Past Year (age 15-24)	288156	.312
Alive	642014	.930	Gave Birth Past Year (age 25-49)	894103	.140
			Ever Married (women age 15-24)	723039	.431
			Ever Married (women age 25-49)	1078875	.936

Notes: All variables, other than log 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 referring to the occurrence of these for the individual in question (if alive) in the preceding two weeks; log weight and log height referring to measurements of living children at the time of the survey.

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 above 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. By including health and family economics, I take

a broader view of consumption than the typical national accounts measure. However, as shown later, this does not drive my results, as these products show close to average growth. 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 for 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.

Before turning to the analysis, it is useful to graphically depict the DHS data that drives the results of this paper. In Figure I below I graph, for each survey x product combination, the country demeaned values of the product consumption[4] against the country demeaned values of the survey year. To provide a money metric for the movements in the consumption of each product, I scale each product measure so that the cross-country standard deviation of the product consumption level equals the cross-country standard deviation of log consumption per equivalent adult reported in the Penn World Tables.[5] Thus, the vertical movement in each product

[4]For the ln variables (rooms, height and weight) I use the urban/rural weighted survey average, whereas for the dichotomous variables I take the logit of that average, i.e. ln(c/(1-c)), as I use the logit as my baseline discrete choice model later in the paper. In each figure I drop the (usually 14) countries for which I have only one survey observation on the product in question. The data of these surveys is used, however, in benchmarking the cross-sectional standard deviation of consumption, as described shortly. I should also note that I drop the middle observation for Nigerian height, as it is bizarrely low and throws off the entire scale of the figure. This observation is used in the analysis below and has little influence, as there are Nigerian surveys before and after it.

[5]Thus, if c_{it} is the country demeaned product consumption measure, c_i the country mean product consumption measure, and σ[PWT] the PWT standard deviation of ln real money consumption levels $\ln(C_i)$ (as reported in Table

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

Durables & Housing Conditions

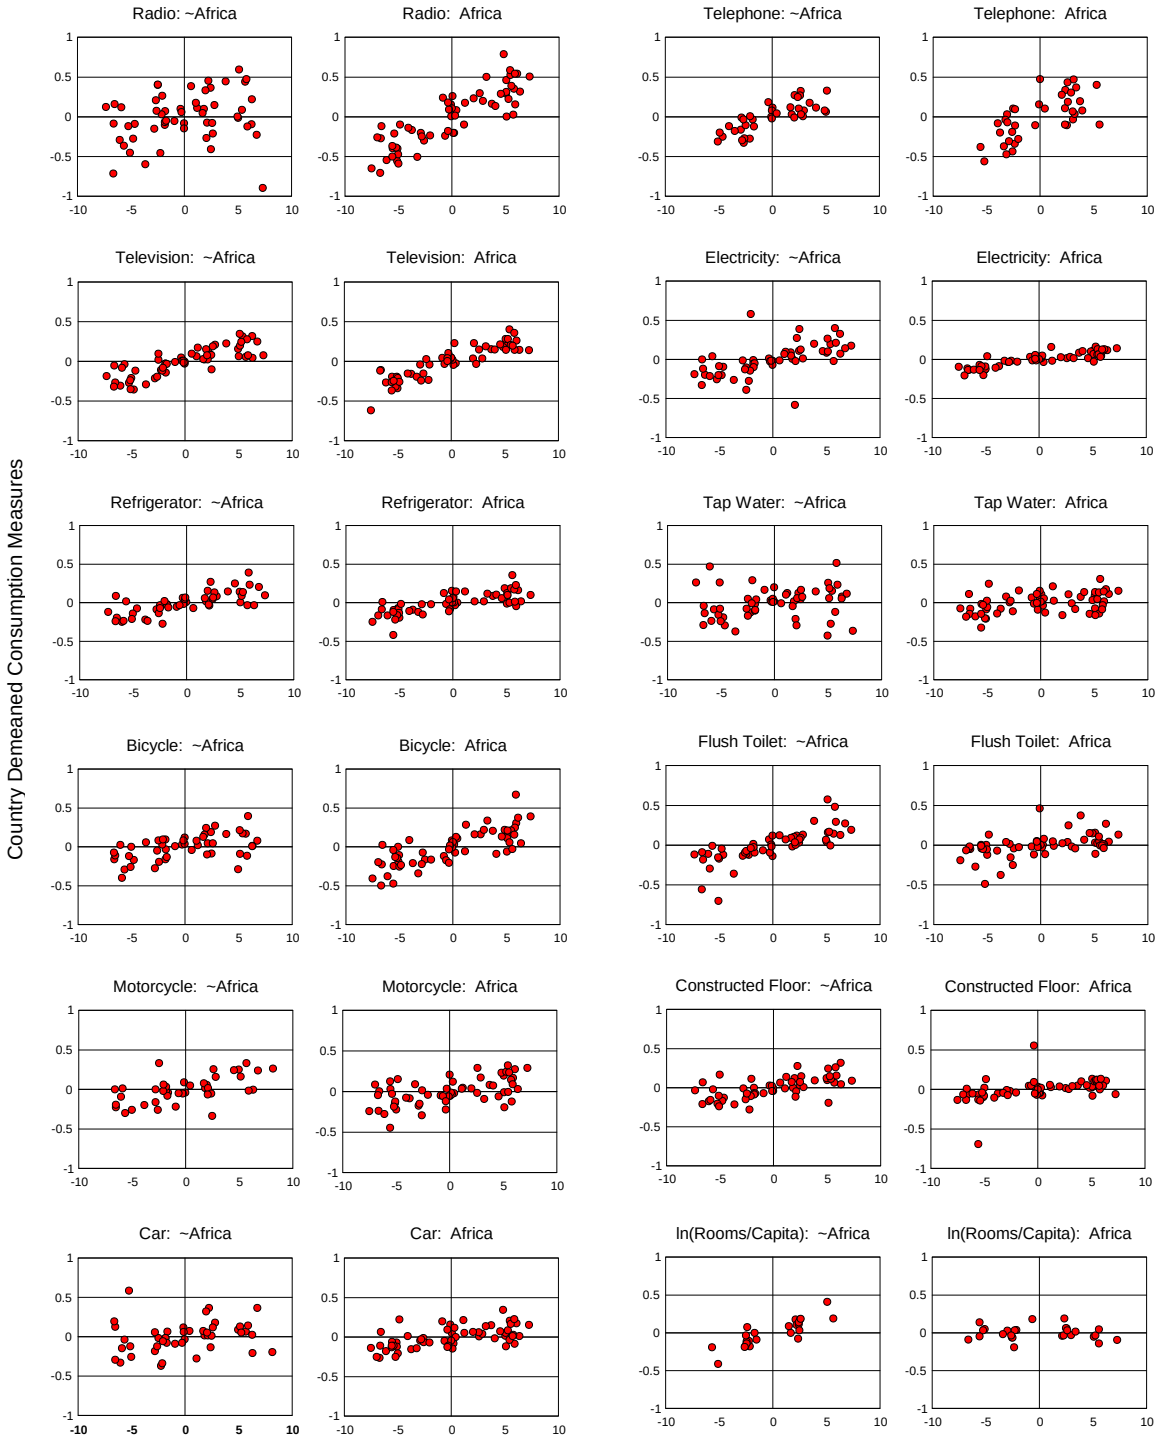

Figure I: Continued

Children's Health

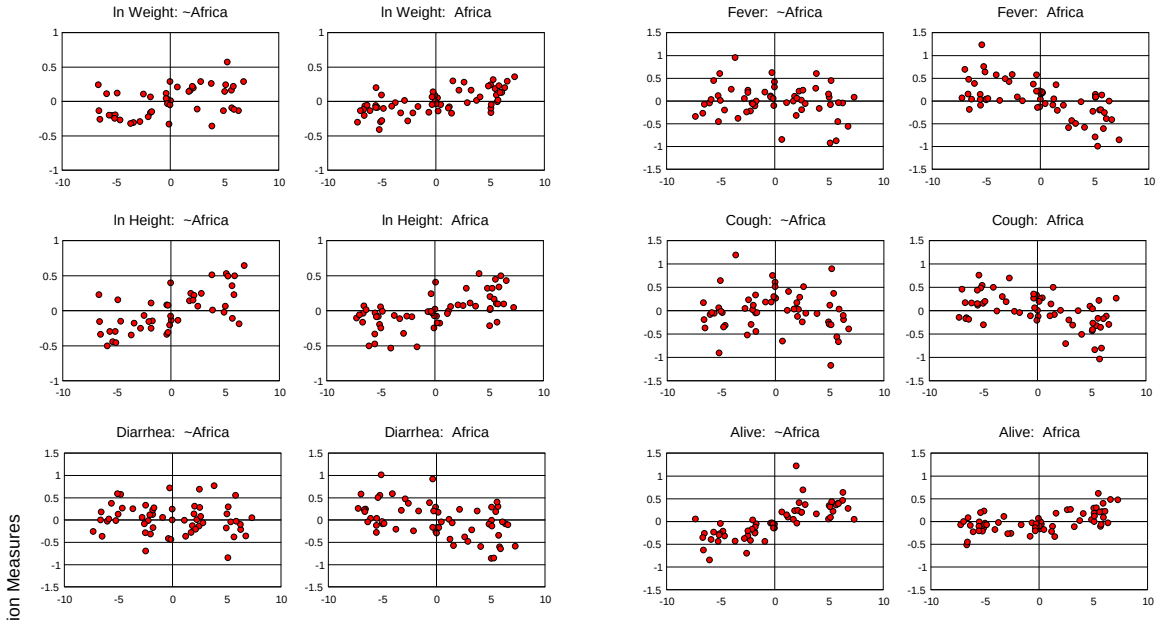

Household Economics

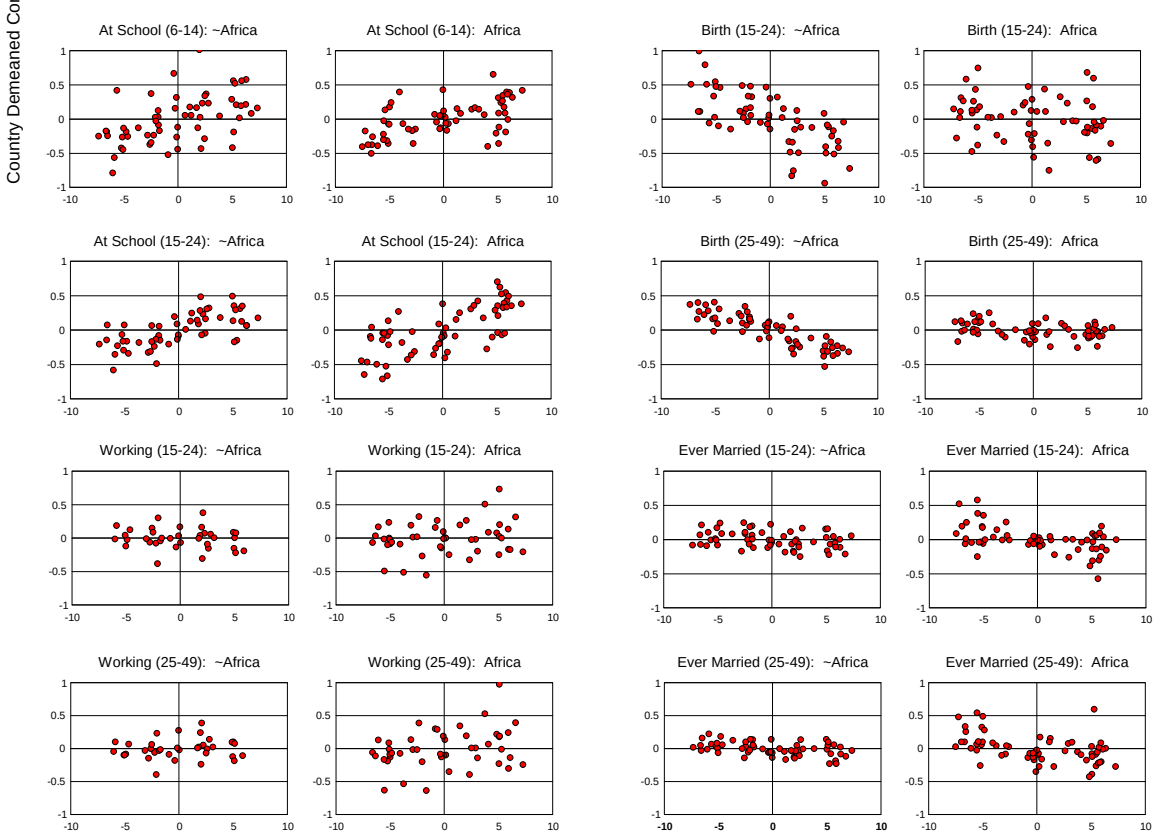

Country Demeaned Consumption Measures

Country Demeaned Year

consumption measure can be interpreted as the money consumption equivalent movement implied by a crude Engel curve calculated off of the cross-national variation in mean product consumption.

The figure shows two characteristics of the DHS data. First, across most products there is simply "too much" movement in consumption, particularly for the African countries. PWT and UN consumption growth rates for sub-Saharan Africa, shown later, are around .01 per annum. Thus, a country (demeaned) year value of 5 in the figure should be associated with a vertical movement of .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.[6] Second, while the PWT and UN suggest that non-African consumption growth is more than double that of sub-Saharan Africa, in the DHS the consumption movement in the African countries appears, by and large, to be roughly equal to that of the non-African countries. A skeptic might argue that my sample of products, however broad I believe it to be, is biased towards a set of goods whose relative prices are falling rapidly, i.e. the LDC equivalent of dvd players in recent decades in the developed world. This, however, cannot explain why African growth in these products matches non-African growth.

III. Methods: An Intuitive Introduction

I begin with an intuitive and simplified presentation of my methods, leaving the more

---

VI later), I divide each c_{it} by $\beta = \sigma[c_i]/\sigma[PWT]$. This can be motivated by the equation $c_i = \beta * \ln(C_i)$. Since this equation should contain an error term, my calculation probably overstates the implied Engel elasticity β and hence understates the growth suggested by the data.

[6]Some products are negatives, e.g. diarrhea, fever and cough, and growth in these cases is defined as a reduction in their incidence in the household. While at this point this may seem arbitrary, in the formal analysis I use the micro-data relationship between the product and educational attainment to determine the change associated with rising consumption. For the reader's information, aside from the three health variables just mentioned, women working when young and births and marriage at any age are found to be negatively associated with household educational attainment (Table V later).

Table II: Average Household Bicycle Ownership and Implied Relative Log Real Consumption in Economies A and B						
	(a) bicycle ownership		(b) equivalent yrs of educ.		(c) log real consumption	
	A	B	A	B	A	B
1990	.220	.200	11.0	10.0	1.10	1.00
2000	.236	.220	11.8	11.0	1.18	1.10
Notes: (a) is the fraction of households owning a bicycle; (b) equals (a) divided by a .02 coefficient derived from a micro data regression of ownership on educational attainment; (c) equals (b) times an estimated .10 Mincerian return to a year of education. All values are hypothetical.						

formal and complete exposition for later. Imagine one observed the data presented in Table II on household ownership of bicycles in two economies. As shown in panel (a), economy A has a higher average ownership level than B and ownership in both economies is growing. Next, consider using micro data in the two economies in both periods to run a regression of ownership on household educational attainment. Say this produces a coefficient of .02 on years of educational attainment. Dividing the mean consumption levels in panel (a) by the coefficient of .02 produces the education equivalent consumption levels reported in panel (b). If one found, separately, that a year's education in both economies results in, say, a 10% increase in log real income and consumption, one could derive the money equivalent log real consumption levels reported in panel (c). We would conclude that economy A was 10% richer than B in 1990, and only 8% richer in 2000, while growth was 8% and 10% in A and B, respectively, between 1990 and 2000. In sum, my approach is to use Engel curves implicitly estimated off of educational attainment data to convert physical consumption levels into money metric measures of real consumption.

Any reasonable reader will immediately object that a host of factors other than real consumption determine the presence of a bicycle in a household. For the purposes of discussion, I will divide these into two categories: (a) influences which increase demand for a given product,

but only at the expense of lowering demand for something else; (b) influences which change measured product demand without reflecting substitution from other products or any changes in underlying real consumption. Relative prices are an obvious cause of the substitution described in (a). Demographic factors contribute to the biases suggested by (b). Thus, households with more members, perhaps in poorer countries or rural areas, are more likely to report the presence of a bicycle for any given level of real living standards per member. Similarly, the height and weight of infants, for any given level of real consumption expenditure, is strongly influenced by their age. I should emphasize that in this characterization of potential problems I exclude factors which lower the overall real price of consumption. Thus, households living in countries where governments provide good transport, power and sanitation infrastructure will, for a given set of nominal goods prices, experience lower shadow prices of consumption and enjoy better measured material outcomes. These should properly be counted as indicative of higher real consumption.

The key characteristic of substitution between products brought about by relative price differences is that it has no particular sign or expected value for any given product. The obvious solution, suggested by sampling theory, is to calculate log consumption values such as those of Table I for a wide variety of products and average these to produce an overall estimate of living standards. To be as representative as possible, the product sample should be "stratified", drawing across diverse areas of expenditure, such as the durables, housing, family economics and health areas indicated in my description of DHS data. Jackknife techniques (i.e. case-wise deletion of observations)[7] and comparison of results across product categories will give a sense of the sensitivity of the results to the product choices.

[7]The application of the jackknife involves calculating a statistic N times, each time deleting one of the N observations. While its principal objective is a non-parametric estimate of the standard error, its calculation allows one to observe and report the sensitivity of the results to individual outliers.

Econometrics provides techniques that improve on the efficiency of simple sample averages. Key among these is the recognition that different observations come with differing degrees of accuracy. Consider, for example, the growth implied by the consumption of a product, as presented in Table II. With $\hat{\beta}_i$ denoting the regression coefficient on educational attainment for product i, M_{itc} its mean consumption level at time t in country c, and R_E the association between log real consumption and education, estimated money metric equivalent growth for product i in country c is given by:

$$(1)\quad \hat{g}_{ic} = R_E \frac{M_{i2000c} - M_{i1990c}}{\hat{\beta}_i} \qquad \hat{\sigma}(\hat{g}_{ic}) = \hat{g}_{ic} \frac{\hat{\sigma}(\hat{\beta}_i)}{\hat{\beta}_i}$$

The right-hand side of (1), the estimated standard error of $\hat{g}_{ic}$, is arrived at through the "delta method" by multiplying the absolute value of the derivative of $\hat{g}_{ic}$ with respect to $\hat{\beta}_i$ by the estimated standard error of $\hat{\beta}_i$.[8] As the equation shows, the standard error of $\hat{g}_{ic}$ will be larger the larger is the ratio of the standard error of $\hat{\beta}_i$ to $\hat{\beta}_i$ itself, i.e. the lower its statistical significance.

Let g_{ic} be the actual Engel curve consumption equivalent growth implied by the growth of the physical consumption of product i. Because of relative price trends, say g_{ic} is distributed normally with mean μ_c (the growth of log real consumption in country c) and variance σ^2. Consequently, an observation $\hat{g}_{ic}$ is normally distributed with mean μ_c and variance $\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2$. Our interest lies in estimating μ_c. The probability or likelihood we observe a sample of N product growth rates for country c is given by:

[8]To keep the example simple, I assume M_{itc} and R_E are known with certainty. In practice, it is the tightness of the Engel curve relation that determines the relative variance of different product observations, as mean consumption levels are estimated to a high degree of accuracy with even modest sample sizes, while R_E affects all products equally.

$$(2)\quad L = \prod_{i=1}^{N} \frac{1}{\sqrt{2\pi\left[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2\right]}} \exp\left[-\frac{1}{2}\frac{[\hat{g}_{ic} - \mu_c]^2}{\left[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2\right]}\right]$$

Taking the derivative of the log of this likelihood with respect to μ_c and setting it equal to zero, we find that the maximum likelihood solution for μ_c is given by:

$$(3)\quad \mu_c = \sum_{i=1}^{N} w_i \hat{g}_{ic} \quad \text{where} \quad w_i = \frac{\left[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2\right]^{-1}}{\sum_i \left[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2\right]^{-1}} \quad \text{and} \quad \sum_i w_i = 1.$$

Thus, under the given distributional assumptions, the most efficient estimate of the growth rate is a weighted average of the estimated product growth rates. The weight placed on each product is declining in its estimated variance. If each product is estimated with the same variance, the weights are all 1/N and we take the simple average across products.[9]

A standard calculation of consumption growth, based upon price and nominal expenditure data, would weight the growth of each product's real consumption by its share of nominal expenditure. Equation (3) shows that, absent such data, my approach uses the significance of the first step estimate of the Engel curve relationship to weight the growth of real consumption implied by dividing product consumption growth by its Engel curve coefficient. In practice, this tends to remove extreme growth outliers as, absent such adjustments, I find African growth to be above 7%, i.e. more than double the 3.4% I report in my variance-adjusted baseline estimates. In addition to accounting for the error with which observations are estimated, I also improve econometric efficiency by introducing run-of-the-mill random effects designed to account for the role relative prices play in producing persistent differences across countries in levels and trends for the consumption of particular products. These also change the relative weighting of

[9]The first order condition for σ is given by $\sum(\hat{g}_{ic} - \mu_c)^2[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2]^{-2} = \sum[\sigma^2 + \hat{\sigma}(\hat{g}_{ic})^2]^{-1}$ which, along with (3), generally gives two non-linear equations in the two unknowns μ_c and σ^2. When each product is estimated with the same variance, this equation has the simple solution $\sigma^2 = \sum(\hat{g}_{ic} - \mu_c)^2 / N - \hat{\sigma}(\hat{g}_{ic})^2$

Table III: Implied Relative Log Real Consumption w/ Adjustment for Demographic Biases								
	(a) dummies		(b) average yrs of educ.		(c) equivalent yrs of education		(d) log real consumption	
	A	B	A	B	A	B	A	B
1990	.140	.150	3.0	2.0	10.0	9.50	1.00	0.95
2000	.150	.130	3.5	4.0	11.0	10.5	1.10	1.05

Notes: (a) reports the dummies in a regression of household ownership on demographic variables, educational attainment and country x time period dummies; (b) equals mean years of household educational attainment; (c) equals (a) divided by the .02 coefficient on educational attainment estimated in (a), plus (b); (d) equals (c) times an estimated .10 Mincerian return to education. All values are hypothetical.

observations, but as they are standard and their empirical influence is trivial, I leave their presentation for later.

Finally, turning to the biases introduced by household demographic characteristics, these can be removed in the micro data regressions. Following on the example earlier above, microdata on household ownership of a bicycle can be run on demographic controls, household educational attainment and a full set of country x time dummies. Say, for the sake of simplicity, that this regression again produces the .02 coefficient on educational attainment described earlier, and the country x time dummies described in panel (a) of Table III. These dummies measure relative consumption purged of the influence of mean demographic variables and educational attainment. My objective is a measure of relative consumption purged only of demographic influences. Consequently, in panel (b) I report the mean household educational attainment in each region x time period which I add to the dummies of panel (a) divided by .02 to produce the regional educational equivalent levels of consumption reported in panel (c). Multiplying these values by the estimate of a 10% income profile of education produces the relative incomes reported in panel (d), which are purged of the confounding influence of demographic factors. The key point of this example is that residual dummy variables from a multivariate regression can

be substituted for mean national consumption levels in calculating the education equivalent consumption levels, thereby correcting for demographic characteristics, provided national mean education levels are added back in, as they are part of the national education equivalent consumption of the product.

Broadly speaking, the type of computations illustrated in Table III, averaged across a variety of products to reduce the error introduced by relative price effects and with the estimation precision and random effects weighting described and alluded to above, form the basis of the calculations central to this paper.[10]

IV. Methods: Product Sampling and the Measurement of Real Consumption

(a) Model

I begin by laying out the theoretical framework and then describe its empirical implementation. Let some measure of the real demand by household h for product p in region r in period t be described by the equation:

$$(4)\ \log(Q_{hprt}) = \alpha_p + \eta_p \log(C^N_{hrt}) + \vec{\xi}'_p \log(\vec{P}_{rt}) + \vec{\beta}'_p \vec{X}_{hrt} + \varepsilon_{hprt}$$

where α_p is a constant, η_p the quasi-income elasticity of demand, C^N_{hrt} nominal household consumption expenditure, $\vec{\xi}'_p$ a vector of own and cross quasi-price elasticities of demand, $\log(\vec{P}_{rt})$ the vector of regional prices relative to some base, $\vec{X}_{hrt}$ and $\vec{\beta}_p$ vectors of demographic characteristics and their associated coefficients, and ε_{hprt} a mean zero idiosyncratic household preference shock. I use the term quasi in describing the elasticities, because

[10] In practice, I calculate urban/rural estimates for each country, and weight these by survey data on the urban/rural household population shares to produce aggregate national estimates of product consumption levels. For the most part, I use discrete choice models rather than linear regressions to calculate regional dummies and educational demand coefficients, so that the estimated household ownership probabilities always lie between 0 and 1. In addition, there is a variant of my procedure where I allow demand patterns to vary country by country (instead of imposing common global patterns), which still allows me to calculate growth rates of real consumption, but not levels. This is explained later in the paper.

$\log(Q_{hprt})$ need not be actual log 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 the quasi-income elasticity of demand equals the negative of the sum of the own & cross quasi-price elasticities:

$$(5)\quad \eta_p = -\sum_q \xi_{pq}$$

Equation (5) 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 (4) in terms of real consumption, we add and subtract from nominal expenditure the expenditure share weighted movement of prices from the base to produce

$$(6)\quad \log(Q_{hprt}) = \alpha_p + \eta_p[\log(C^N_{hrt}) - \vec{\Theta}'_{rt}\log(\vec{P}_{rt})] + \eta_p[\vec{\Theta}'_{rt} + \vec{\xi}'_p/\eta_p]\log(\vec{P}_{rt}) + \vec{\beta}'_p\vec{X}_{hrt} + \varepsilon_{hprt}$$

where $\vec{\Theta}_{rt}$ is a vector of regional product expenditure shares.[11] The second term on the right-hand side is real expenditure, while the third term can be thought of as a region x time error term:

$$(7)\quad \log(Q_{hprt}) = \alpha_p + \eta_p\log(C^R_{hrt}) + \eta_p\varepsilon^{\vec{P}}_{prt} + \vec{\beta}'_p\vec{X}_{hrt} + \varepsilon_{hprt}$$

where the superscript $\vec{P}$ on $\varepsilon^{\vec{P}}_{prt}$ is used to emphasize the role relative prices play in determining this 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 regional error term which, when normalized by the quasi-income elasticity, is a zero-

[11] 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

weight average of relative price changes; something that, arguably, is homoskedastic across products and has an expected value of zero.

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:

$$(8)\quad \log(C^R_{hrt}) = \alpha_{rt} + \log(Y^R_{hrt}) \qquad \log(Y^R_{hrt}) = \log(Y^{R\sim E}_{rt}) + R_E E_{hrt}$$

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 $\log(Y^{R\sim E}_{rt})$ is education adjusted log regional real income at time t. It follows that average regional log household consumption expenditure at time t is given by:

$$(9)\quad \log(C^R_{rt}) = \log(C^{R\sim E}_{rt}) + R_E E_{rt}$$

where E_{rt} is mean regional household educational attainment and $\log(C^{R\sim E}_{rt}) = \alpha_{rt} + \log(Y^{R\sim E}_{rt})$ is education adjusted log regional real expenditure per adult.[12] Average log country expenditure is the population weighted sum of log regional real expenditure:

$$(10)\quad \log(C^R_{ct}) = \sum_{r\in r(c)} S_{rt} \log(C^R_{rt})$$

where *r(c)* is the set of regions in country *c* and the S_{rt} are the regional population shares. Regions 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

Finally, I assume that real consumption expenditure is growing at an average rate g, so that real household consumption in country c at time t can be written as:

$$(11)\quad \log(C^R_{ct}) = \log(C^R_c) + gt + g_c t + \varepsilon_{ct}$$

[12]Clearly, savings rates are allowed to vary across regions and time (note α_{rt} in (8)), but there is the implicit assumption that savings rates out of permanent income do not vary by educational attainment. This allows me to estimate the relative real consumption expenditure of educational categories using data on their relative incomes.

where g_c represents the deviation of the country's growth rate from the average g and $\log(C_c^R)$ equals log relative consumption in the base year, which in my analysis will be the year 2000. Uncovering the base year levels $\log(C_c^R)$ and average growth rate g of real country log consumption is the fundamental objective of my analysis.

(b) Estimation

Estimation proceeds in two steps. In the first step, I combine all of my surveys to estimate household demand equations, product by product, of the form

$$(12)\quad \log(Q_{hprt}) = a_{prt} + b_p E_{hrt} + \vec{c}_p' \vec{X}_{hrt} + e_{hprt}$$

where $\log(Q_{hprt})$ will usually be the index in a discrete choice probability model or otherwise the log of some measurable continuous outcome, and where the a_{prt}s are a complete set of product specific region x time (equivalently, survey[13]) dummies. Under the assumptions laid out above, asymptotically the coefficient estimates converge to the following values:

$$(13)\quad \hat{b}_p = \eta_p R_E \qquad \hat{\vec{c}}_p = \vec{\beta}_p \qquad \hat{a}_{prt} = \alpha_p + \eta_p \log(C_{rt}^{R\sim E}) + \eta_p \varepsilon_{prt}^{\vec{P}}$$

While the unconditional expectation of $\varepsilon_{prt}^{\vec{P}}$, 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 dummies.

Next, I construct measures of log real regional consumption as implied by the consumption of a particular product by dividing the product x region x time dummy by the coefficient on educational attainment, adding the survey estimate of average regional educational attainment, and multiplying by a separately estimated return to education:

[13]In practice, I assign a common date (equal to the mean household survey date) to all observations within a particular country survey. Thus, the t's in the equation above are really country survey dates.

$$(14)\quad \log(\hat{C}^R_{prt}) = \hat{R}_E\left[\frac{\hat{a}_{prt}}{\hat{b}_p} + \hat{E}_{rt}\right]$$

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

$$(15)\quad \log(\hat{C}^R_{pct}) = \sum_{r\in r(c)} S_{rt}\log(\hat{C}^R_{prt})$$

These estimates are then projected on product and country dummies, time entered separately for the sub-Saharan African and non sub-Saharan African countries, and a series of random shocks designed to improve econometric efficiency:

$$(16)\quad \log(\hat{C}^R_{pct}) = a_p + a_c + g_A t_A + g_{\sim A} t_{\sim A} + v_c t + v_p t + u_{pc} + e_{pct} + \hat{e}_{pct}$$

In (16), having removed variation in mean product consumption levels with the product constants a_p,[14] I use a_c to estimate $\log(C^R_c)$, the relative country consumption level in the base year (2000), and g_A and $g_{\sim A}$ to estimate the mean African and non-African consumption growth rates. The random coefficients v_c and v_p explicitly allow growth to vary across countries and, due to relative price trends, across product types, while the random effect u_{pc} takes into account the fact that relative price differences will result in persistent differences in product consumption levels across countries. Each random shock is independently drawn at the level of its subscript(s). Thus, v_c is an independent draw from a zero-mean normal distribution affecting the growth of country c, while u_{pc} is an independent draw from a zero-mean normal distribution affecting the level of consumption of product p in country c The regression residual variation has two components:

[14]This is unnecessary for a balanced panel, but is important for unbalanced panels as otherwise mean worldwide product consumption levels have a spurious influence on the estimates of relative country aggregate consumption.

(a) the residual variation of the true $\log(C^R_{pct})$ after accounting for the components modeled on the right-hand side, e_{pct}; plus (b) the additional variation introduced by the use of the estimate $\log(\hat{C}^R_{pct})$ of $\log(C^R_{pct})$ as the dependent variable, $\hat{e}_{pct}$.

By explicitly stating the likelihood I can provide the reader with a fuller description of the role played by the different components in (16). Under the assumption that all of the errors and random shocks are normally distributed, the probability the sample is observed is given by

$$(17)\quad L = \frac{\exp[-.5(Y - X\beta)'\Omega^{-1}(Y - X\beta)]}{(2\pi)^{N/2}|\Omega|^{1/2}} \quad \text{where} \quad \Omega = \Sigma(\text{RS}) + \text{I} \times \sigma[e_{pct}]^2 + \hat{\Sigma}(\text{FS})$$

and where Y is the Nx1 vector of observations $\log(\hat{C}^R_{pct})$, X is the Nxk matrix of regressors consisting of product and country indicator variables and time entered separately for the African and non-African countries, and β is kx1 made up of the coefficient vectors a_p and a_c, plus g_A and $g_{\sim A}$. The covariance matrix Ω is made up of three components: (1) Σ(RS), the covariance across observations created by the random shock $v_c t + v_p t + u_{pc}$, which will depend on the standard deviations of the component processes, $\sigma[v_c]$, $\sigma[v_p]$ and $\sigma[u_{pc}]$; (2) $\text{I}\times\sigma[e_{pct}]^2$, the orthogonal variation stemming from the residual orthogonal variation in $\log(C^R_{pct})$; and (3) Σ(FS), the covariance across observations stemming from the covariance in the estimation error $\log(\hat{C}^R_{pct}) - \log(C^R_{pct})$. The log likelihood is maximized with respect to β, $\sigma[v_c]$, $\sigma[v_p]$, $\sigma[u_{pc}]$ and $\sigma[e_{pct}]$. Σ(FS) is fixed and is calculated from the first step covariance matrices.[15]

[15] As shown in (15) and (14), $\log(\hat{C}^R_{pct})$ 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. Thus, in precisely the cases where I want to place the least weight on a variable (because the estimate of b_p has a substantial probability mass around zero), the delta method will be a poor guide to Σ(FS). I handle this problem by using Monte Carlo techniques to estimate Σ(FS), generating 100,000 draws from the estimated joint distribution of the a_{prt}s and b_p in each product

Maximization of (17) with respect to β produces the standard GLS estimate of the coefficient vector as a weighted average of the X and Y observations:

$$(18)\quad \hat{\beta} = (X'\Omega^{-1}X)^{-1}X'\Omega^{-1}Y$$

In this case, the weighting has two components. First, there is the weighting imposed by the random shocks. Thus, for example, to the extent $\sigma[v_c]$ and $\sigma[v_p]$ are found to be large, in estimating g_A and $g_{\sim A}$ less than one-for-one weight will be placed on countries or products with relatively large numbers of time series observations, reflecting the fact that, because of the covariance of growth within countries or products, large samples for a given country or product provide less information than equivalent samples drawn across countries or products. Similarly, if $\sigma[u_{pc}]$ is found to be large, less than one-for-one weight will be placed on large numbers of product x country observations in estimating the product and country means a_p and a_c.[16] Given the highly unbalanced nature of my panel, these adjustments could have a large effect on the coefficient estimates if there is a great deal of variation in growth rates and levels by sub-sample size. In practice, they do not, as shown further below.

The second component of weighting in (18) involves the covariance matrix of the first step estimates of $\log(\hat{C}^R_{pct})$, Σ(FS). If one orders the observations product by product, one sees that this covariance matrix is largely[17] block diagonal, made up of the product specific matrices Σ_p(FS). The inverse of a block diagonal matrix is itself block diagonal. Thus, *Y-XB* deviations

---

equation and then calculating the resulting mean and variance of the ratios, to which I then add the covariance matrix of the estimated mean educational attainment by region.

[16] Appendix B provides some mathematical exposition to back up this intuition.

[17] Since the components a_{pct} and b_p are estimated product by product (i.e. independent variables are entered separately for each product), the maximum likelihood estimate of their covariance matrix is block diagonal. $Log(\hat{C}^R_{pct})$ also depends on the estimate of mean regional attainment E_{rt}, which is common to all products. However, the estimated variance of E_{rt} is tiny relative to the product specific components.

for products where the first step covariance matrices are large will face small inverses, placing correspondingly small weight on those observations.[18] $\log(\hat{C}^{R}_{pct})$ depends on the ratio of the regional dummies a_{pct} to the education Engel coefficients b_p (see equation 14). Since, absent other regression components, regional dummies are generally estimated quite accurately in large samples, this covariance matrix is large primarily when the consumption-education relation is weak. Thus, as in the case of the simple example of the previous section,[19] my estimates place more weight on products where the estimated relationship between education and consumption is stronger. As shown further below, this weighting is extremely important as absent this adjustment average growth rates are found to be 4.7 and 7.2 percent for the non-African and African economies, respectively.

Finally, I should note that when comparing individual country levels to PWT levels, I estimate the country levels a_c as fixed effects, as described above. However, the standard deviation of a set of point estimates is inflated by estimation error. Consequently, when I seek to describe the standard deviation of country levels to compare with the same statistic from PWT, I estimate the country levels as random effects u_c with standard deviation $\sigma[u_c]$. This choice of specification has a negligible effect on the other coefficients estimated in the regression.

[18]The reader will recognize that for heuristic purposes I am acting as if $[\Sigma(RS)+I\sigma[e_{pct}]^2+\Sigma(FS)]^{-1} = \Sigma(RS)^{-1}+I/\sigma[e_{pct}]^2+\Sigma(FS)^{-1}$. This is, of course, not true, so the description in the text literally applies only when the other components are removed from the model, i.e. ignoring interactions between the component matrices. In the case of this paper, the presence of $\Sigma(RS)$ does not really affect the estimates so the interactions between the random shocks and estimation accuracy weighting are, indeed, unimportant.

[19]In that simple example, I focused on the calculation of a mean across product observations for a single country. With each product estimated separately, that produced a diagonal covariance matrix, allowing me to use simple algebra to discuss the individual product likelihoods. My actual estimates involve calculations for groups of countries in an unbalanced panel, combined with random shocks across products and countries, all of which produces the more complicated matrix algebra discussed above. The intuition, however, is the same.

V. Results: 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 A). I run the typical Mincerian regression of log wages on educational attainment, age, sex and regional controls.

As shown in Table IV, 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 geographical 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 centers) which provide higher earnings for any given level of education. As more detailed geographical 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 is also important to note that more detailed geographical controls increase the noise to signal ratio in educational attainment, biasing the coefficient towards zero. This is particularly relevant for

Table IV: Log Wage Regressions					
	(1) survey dummies	(2) survey x rural/ urban dummies	(3) cluster random effects	(4) cluster fixed effects	(5) cluster fixed effects (IV)
educ	.115 (.002)	.108 (.001)	.104 (.001)	.095 (.002)	.116 (.005)
age	.047 (.007)	.047 (.007)	.049 (.006)	.048 (.007)	.046 (.008)
age2	-.000 (.000)	-.000 (.000)	-.000 (.000)	-.000 (.000)	-.000 (.000)
sex	-.350 (.019)	-.360 (.019)	-.365 (.015)	-.366 (.017)	-.396 (.020)
N	22996	22996	22996	22996	18418
Notes: Dependent variable is log annualized work income of individuals aged 25-65 working for others. Coefficients on age2 are small (around -.0004) but significant. Educ. and age measured in years; sex = 1 if female.					

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 IV controls for 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.[20] 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.[21] This implies a measurement standard error of about 1.6, i.e. that about 32 percent of the wage reporting sample, with mean educational attainment of 9.5 years[22] over or understate their educational attainment

[20]The absolute values of the t-statistics of these four variables in the first stage regression are 45.1, 4.1, 5.7 and 6.1, respectively.

[21]The education coefficient of column (4) using the sample of column (5) is 9.40. Divided by column (5)'s coefficient of 11.60 this indicates a signal to signal plus noise ratio of .81. The measurement standard error reported in the next sentence equals the square root of .19 times the variation in education orthogonal to the other controls.

[22]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 (5.1 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 5.4 years, while the average educational attainment of adults who report earnings data, whether working for themselves or others, is 6.7 years. If I rerun the specification of column (5) using all adult individuals reporting earnings from work (including, presumably, capital income) I get an education coefficient of 13.6. Thus, a

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 an "attenuation bias adjusted" return to education of 12.5 percent in that column. When compared with column (5)'s point estimate, this indicates that although 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 (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.[23] Readers who have strong alternative priors can scale all of the growth rates and cross-national standard deviations of real expenditure reported below by the ratio of their preferred number to 11.6. However, it would take an enormous reduction in the estimated return to education, to around 3 percent, to bring the DHS-implied African growth figures in line with international estimates.

---

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. However, 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, just to report what the standard selectivity adjustments produce, 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 12.0 in the second.

[23]In a later section I allow R_E to vary by region and find that it is systematically higher in sub-Saharan Africa, which raises the estimated growth for that region.

Table V: Product Level Estimates of the Response to Educational Attainment					
	coefficient	Y elasticity		coefficient	Y elasticity
Radio	.153 (.001)	0.57	Electricity	.228 (.001)	0.92
Television	.236 (.001)	1.21	Tap Drinking Water	.076 (.001)	0.36
Refrigerator	.253 (.001)	1.64	Flush Toilet	.234 (.001)	1.37
Bicycle	.056 (.001)	0.34	Constructed Floor	.210 (.001)	0.73
Motorcycle	.190 (.001)	1.47	log(Rooms/Capita)	.020 (.000)	0.17
Car	.250 (.001)	2.01			
Telephone	.248 (.001)	1.77	At School (6-14)	.200 (.001)	0.50
			At School (15-24)	.148 (.001)	0.84
log Weight	.007 (.000)	0.06	Working (15-24)	-.032 (.002)	-0.16
log Height	.002 (.000)	0.02	Working (25-49)	.020 (.001)	0.08
Diarrhea	-.033 (.001)	-0.23	Birth (15-24)	-.012 (.001)	-0.07
Fever	-.019 (.001)	-0.11	Birth (25-49)	-.026 (.001)	-0.19
Cough	-.006 (.001)	-0.04	Marriage (15-24)	-.058 (.001)	-0.28
Alive	.059 (.002)	0.04	Marriage (25-49)	-.077 (.001)	-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: log number of persons in the household; (2) children's health: sex, log(1+age in months) and log(1+age in months) squared (for all but height, weight and mortality, which are quite linear in log(1+age)); (3) household economics: age and age squared, as well as sex for education attendance variables (all others refer to women alone). Y elasticity is the income elasticity, as explained in the footnote in the text. Each equation is estimated separately.

Moreover, such a reduction would produce new puzzles, as it would imply very low growth outside of Africa and an extremely small cross-country variation in living standards.

(b) First Step Estimates

Table V above 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 the household and individual demographic controls noted in the table. With the exception of log weight, height and rooms per capita, the figures are the coefficients in a logit discrete choice model with the implied quasi-income elasticities evaluated at the sample mean probability.[24]

[24]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 $\beta(1-P)/R_E$, where P is the mean sample value (Table I).

For our purposes, the main relevance of Table V is that it establishes that each of the real consumption variables used in this paper is significantly and substantially related to real income, as measured by years of education. Across the different products, none of the coefficients is even close to being insignificant at the 1% level. The income elasticities, coupled with the standard deviation (s.d.) of mean household adult education (4.5 years) and implied s.d. of predicted log incomes (4.5*.116 ≈ .5), produce substantial variation in predicted outcomes. Thus, a one s.d. movement in educational attainment produces a log 28 percent higher relative probability of owning a radio (mean value of .573 - see Table I) and a log 68 percent higher probability of having a flush toilet (.323). 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 42 percent (.340) with a s.d. movement in incomes, while the probability a young woman is working (.412) or ever-married (.431) falls by 8 percent and 14 percent, respectively.

(c) The Growth and Standard Deviation of Real Consumption

Table VI below estimates the growth and standard deviation of living standards in my sample of African and non-African countries. I begin by establishing, as a benchmark, the Penn World Tables & United Nations national accounts measures of consumption growth and relative levels.[25] The two data sources are broadly in agreement, suggesting a non-African growth rate of

[25]To make the results comparable with what follows, these estimates are based upon the 135 country x year combinations present in my 1990-2006 DHS data.

Table VI: Estimates of the Growth and Standard Deviation of Living Standards					
PWT & UN national accounts: $y_{ct} = a + g_{\sim A}*t + g_A*t + u_c + v_c*t + e_{ct}$					
	Penn World Tables 7.0 Private Consumption			UN National Accounts Private Consumption	
	per Capita		per Equivalent Adult	per Capita	
$g_{\sim A}$	.022 (.004)		.020 (.004)	.022 (.004)	
g_A	.011 (.003)		.011 (.003)	.009 (.003)	
$\sigma[u_c]$	.818 (.078)		.790 (.075)	.710 (.068)	
$\sigma[v_c]$	.010 (.003)		.010 (.003)	.011 (.003)	
$\sigma[e_{ct}]$	.084 (.010)		.083 (.009)	.080 (.009)	
DHS products: $y_{pct} = a_p + g_{\sim A}*t + g_A*t + u_c + v_p*t + v_c*t + u_{pc} + e_{pct}$					
	All Products	Consumer Durables	Housing	Health	Family Economics
$g_{\sim A}$	.038 (.006)	.046 (.010)	.038 (.011)	.033 (.006)	.031 (.006)
g_A	.034 (.005)	.056 (.010)	.018 (.011)	.034 (.006)	.025 (.006)
$\sigma[u_c]$	.713 (.072)	.742 (.090)	1.08 (.123)	.578 (.068)	.592 (.071)
$\sigma[v_p]$	.019 (.003)	.024 (.007)	.017 (.006)	.006 (.005)	.010 (.005)
$\sigma[v_c]$	.015 (.002)	.016 (.004)	.027 (.005)	.013 (.005)	.013 (.003)
$\sigma[u_{pc}]$	.872 (.020)	.968 (.042)	1.01 (.053)	.504 (.030)	.765 (.036)
$\sigma[e_{pct}]$	.241 (.006)	.221 (.009)	.252 (.014)	.273 (.018)	.206 (.010)

Notes: The u terms represent random effects allowing for variation in country and country x product levels, the v terms represent random variation in country and product growth rates, and e represents the error term. The subscripts denote the index across which the random shock or error applies (e.g. v_c is random variation in country growth). The PWT and UN regressions do not include random product level and growth variation because the dependent variable is a national GDP aggregate. $\sigma[.]$ represents the estimated standard deviation of the relevant random effect or error. PWT uses PPP measures of real consumption and the UN measures are in constant market exchange US dollars with ad hoc PPP adjustments (see footnote in text). PWT calculates equivalent adults by assigning a weight of .5 to persons under 15. DHS measures incorporate the first step covariance matrix into the likelihood, as discussed earlier above.

just over 2 percent, a sub-Saharan growth rate of around 1 percent, and a standard deviation of living standards across countries in 2000 (the base year) of between .7 and .8.[26] As shown in the lower left hand panel of the table, the DHS product data are consistent with a comparable

[26]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 (as reported at http://unstats.un.org/unsd/snaama/formulas.asp, 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").

standard deviation of living standards in 2000 (.713), but suggest a non-African growth rate of 3.8 percent and a sub-Saharan growth rate of 3.4 percent, the latter being three and a half times that reported by the PWT and UN. When the DHS data are examined product group by product group, we find greater sub-Saharan growth in durable goods (5.6 percent) and lower growth in housing (1.8 percent), but even this measure is still double that of the international sources. The consumption growth implied by health and family economics is slightly below the average for all product groups. Hence, my results do not stem from the fact that I use a concept of consumption that is broader than the typical national accounts measure.[27] Finally, I note that the standard deviation of living standards is substantially higher in housing, but the overall dispersion of these measures by product group is not grossly inconsistent with the PWT aggregates.

Figure II graphs the DHS point estimates of relative consumption levels in 2000 (the base year) against the comparable estimates from the PWT. For the purposes of comparison, I show data from PWT 6.2, the earliest to contain 2000 data for all my economies, and the latest PWT 7.0, which incorporates significant updates based upon the 2005 ICP worldwide detailed study of prices. Several facts stand out. First, the most recent version of PWT contains a massive downward revision of the relative consumption of Zimbabwe, producing a huge discrepancy with my DHS estimate. In a hyperinflationary economy small differences in the timing of the measurement of nominal expenditure and price levels can produce extraordinary errors, and I would be inclined to favour my DHS estimates or, if necessary, the earlier PWT calculations. Second, my DHS estimates are systematically higher than the PWT for the former centrally planned economies which, because the material product system didn't measure non-material

[27]Restricting my measure to durables and housing together, I get non-African and African growth rates of 4.3 (.009) and 4.1 (.009) percent, respectively.

Figure II: Relative Real Consumption (2000)

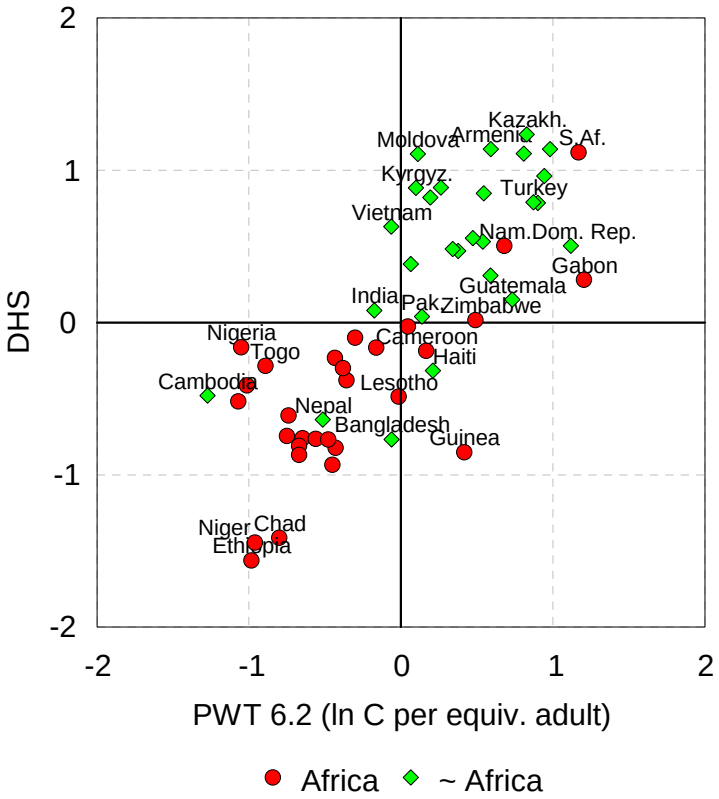

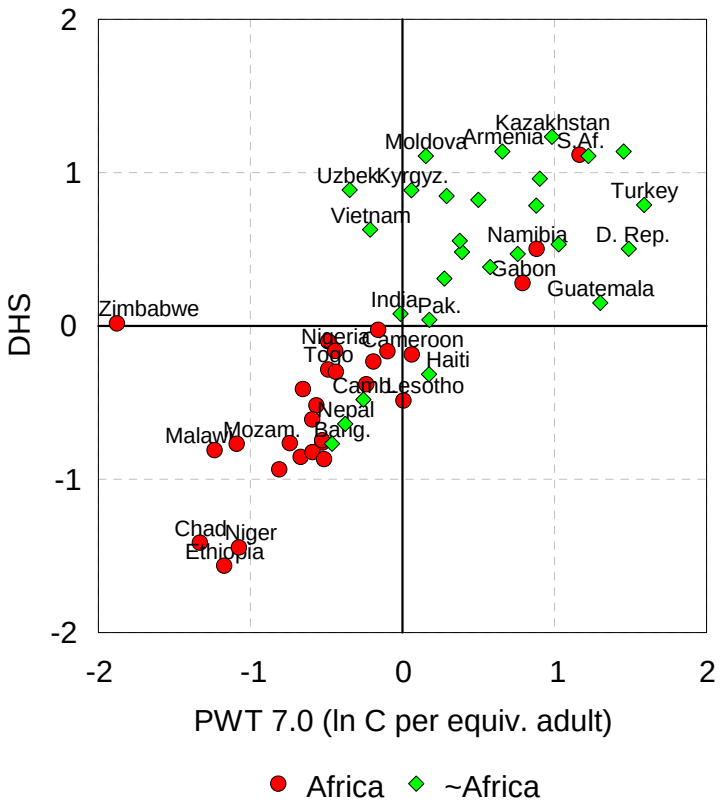

sectors such as services, tend to underestimate GDP.[28] Excluding Zimbabwe and the former centrally planned economies, the correlation between the DHS and PWT 7.0 relative level estimates for the year 2000 is .902. In PWT 7.0 the sub-Saharan economies are on average 97% poorer than the non-African countries. My DHS estimates return a similar log gap of .98.[29]

Figure II illustrates a third significant fact. Between PWT 6.2 and PWT 7.0 there is a strong convergence towards my DHS calculations, as evidenced by the tighter fit around a 45^{o} line in the second panel. Much of this stems from the fact that no benchmark study of prices existed for many of the countries in PWT 6.2. Regressing the change in the estimate of relative consumption between PWT 6.2 and PWT 7.0 on the difference between the PWT 6.2 and my DHS estimates of relative consumption, I get a coefficient of -.47 (-.39 without Zimbabwe). If, however, I restrict attention to the 16 economies for which no benchmark study of prices existed in PWT 6.2 (which does not include Zimbabwe), I get a coefficient of -.66. As the PWT has developed actual data on prices for some of its economies, and improved the estimates of the others with the detailed 2005 ICP, its estimates of living standards in the year 2000 have converged to those I derive from the DHS. The PWT estimates of growth in the poorest regions of the world, however, remain dependent upon the largely fabricated historical series of GDP

[28]Thus, in the case of China, the example I am most familiar with, as surveys have been initiated to cover previously unmeasured sectors there have been large upward revisions of GDP. I should also note that this discrepancy is not due to my use of non-traditional consumption measures such as health and family economics. The average gap between the DHS and PWT estimates of the relative GDP of the 7 former CPEs in Figure I is 62%. If I recalculate the DHS estimates without health and family economics, it actually rises to 71%.

[29]Some readers have queried whether this, coupled with my estimates of African growth, does not imply implausible poverty in Africa prior to the base year 2000. In response, I ask that the following facts be kept in mind: (1) the gap between the highest and lowest log country consumption per equivalent adult in 2000 in PWT 7.0 is 5.0, or 3.5 if restricted to the 56 countries I study; (2) PWT 6.2 showed a log gap of .69 in the base year, thus the PWT revision alone moved relative African incomes down by almost 30 percent; (3) My analysis is for 1990-2006, so all I am arguing is that rather than losing 1% per annum from 1990 to 2000 relative to the other LDCs in my sample (as suggested by PWT & UN) Africa kept pace with them; (4) In an absolute sense I am reporting .34 growth for Africa from 1990 to 2000 as opposed to the .11 indicated by PWT 7.0. In sum, compared to the differences within and across versions of PWT the relative and absolute movements I am talking about are quite small.

Table VII: Crude Growth of Living Standards by Product (regression of country x product measures on trends and country dummies)					
	$g_{\sim A}$	g_A		$g_{\sim A}$	g_A
Radio	.016 (.008)	.056 (.007)	Electricity	.056 (.008)	.048 (.007)
Television	.055 (.006)	.067 (.006)	Tap Drinking Water	.008 (.022)	.028 (.020)
Refrigerator	.040 (.006)	.029 (.005)	Flush Toilet	.068 (.010)	.019 (.009)
Bicycle	.082 (.019)	.131 (.015)	Constructed Floor	.032 (.007)	.019 (.006)
Motorcycle	.035 (.008)	.027 (.006)	log(Rooms/Capita)	.040 (.013)	-.015 (.010)
Car	.016 (.006)	.016 (.005)			
Telephone	.081 (.016)	.081 (.016)	At School (6-14)	.034 (.007)	.044 (.006)
			At School (15-24)	.035 (.009)	.028 (.008)
log Weight	.027 (.010)	.032 (.008)	Working (15-24)	.027 (.067)	-.046 (.049)
log Height	.055 (.043)	.019 (.034)	Working (25-49)	.029 (.113)	.156 (.082)
No Diarrhea	.016 (.028)	.076 (.025)	Birth (15-24)	.149 (.029)	.038 (.026)
No Fever	.048 (.056)	.245 (.049)	Birth (25-49)	.118 (.014)	.021 (.013)
No Cough	.105 (.193)	.542 (.170)	Marriage (15-24)	.026 (.011)	.050 (.009)
Alive	.083 (.010)	.039 (.009)	Marriage (25-49)	.027 (.010)	.046 (.009)
Product Averages: $g_{\sim A}$ = .050 (.010), g_A = .069 (.008).					
Note: Dependent variable in each case is the product x country level given by equation (15).					

growth circulated by international agencies.[30]

Since the key discrepancy between my results and international sources lies in the growth rate, in Table VII I summarize this aspect of the DHS data by reporting the consumption growth estimated by simply regressing the real consumption levels implied by each DHS product (see equation (15) earlier) on time trends and country dummies. These numbers highlight two aspects of my results and methodology. First, the average sub-Saharan product growth rate, at 6.9 percent, is higher than the average non-African product growth rate of 5.0 percent, suggesting

[30]There has been a slight upward revision of growth rates between PWT 6.2 and PWT 7.0, as the analysis of the upper half of Table VI produces slightly lower growth rates using PWT 6.2 data (e.g. growth of 1.7 percent outside of sub-Saharan Africa and 0.9 percent within sub-Saharan Africa using consumption per equivalent adult). This should represent revision of national accounts measures, and not PWT PPPs, as the PWT measures of GDP by component (e.g. consumption) simply involve extrapolating levels in the benchmark year using national accounts growth rates. Thus the inconsistency in PWT growth rates produced by the reweighting of GDP components in each new benchmark highlighted by Johnson et al (2009) is not relevant here.

Table VIII: Sensitivity Tests

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

	first step logit for dichotomous variables					
	2nd step w/out Σ(FS)	2nd step w/out Σ(RS)	Jackknife Products	Bootstrap All Steps	1st step cluster random effects	1st step cluster fixed effects
$g_{\sim A}$	.047 (.019)	.035 (.006)	.038 (.005)	.038 (.008)	.047 (.006)	.049 (.008)
g_A	.072 (.018)	.032 (.005)	.034 (.005)	.036 (.008)	.038 (.006)	.038 (.007)
$\sigma[u_c]$	.938 (.117)	.743 (.073)	.713 (.083)	.739 (.092)	.841 (.085)	.853 (.087)
	alternative first step functional forms					
	Probit	Weibull	Gompertz	Cauchy	Linear	Hermite
$g_{\sim A}$	.037 (.005)	.039 (.005)	.041 (.007)	.046 (.007)	.037 (.005)	.038 (.005)
g_A	.032 (.005)	.028 (.005)	.042 (.006)	.041 (.007)	.029 (.005)	.032 (.005)
$\sigma[u_c]$	.680 (.069)	.675 (.069)	.820 (.083)	.957 (.102)	.657 (.067)	.692 (.070)

Note: Unless otherwise noted, each specification includes the full set of error terms (v_p, v_c, u_{pc}, e_{pct}) as in the lower left panel of Table VI, but only the g_i & $\sigma[u_c]$ are reported. "w/out Σ(FS)": without first step estimation error covariance matrix in 2nd step GLS covariance matrix. "w/out Σ(RS)": without covariance matrix induced by random shocks v_p, v_c and u_{pc} in 2nd step GLS covariance matrix (includes the random effect u_c as this is used to measure dispersion of base year consumption levels).

that overall African consumption growth is at least on par with non-African growth.[31] Second, these numbers show that, in producing the estimates of Table VI above my method of weighting by including the first step covariance matrix in the GLS likelihood systematically places a lower weight on high growth outliers. This is further emphasized in the upper left hand panel of Table VIII below, where I calculate the aggregate consumption growth implied by the DHS data using the same random effects model specified in Table VI, but without the inclusion of the first-step covariance matrix in the likelihood. In this (econometrically incorrect) specification, I find average growth rates of 4.7 percent and 7.2 percent in the non-African and sub-Saharan countries, respectively, and a much higher cross country standard deviation of .938 in the year 2000.

[31]For the reader who notes it, I should explain that the large negative growth implied by the market participation of young women comes from the fact that in the micro-data regression young women's participation is negatively associated with household educational attainment (i.e. young women in richer households are less likely to be working and more likely to be in school), but the trend in the African sample is for rising market participation by young women. However, neither the African nor the non-African trend in this regression is significant.

Beyond estimation without the covariance matrix, Table VIII reports additional sensitivity tests of the DHS results. In the second upper-left column, I estimate the baseline model without the random effects for country x product consumption levels (u_{pc}) and without the random variation in product and country growth rates (v_p and v_c). Relative to this panel, we see that the baseline model (lower-left panel of Table VI) has slightly higher growth rates. As noted earlier, the controls for random variation in product and country growth rates (v_p and v_c) reduce the relative weight on products or countries with large numbers of observations, which could be important in my unbalanced panel. Although the estimated standard deviations of these shocks in the baseline model are quite substantial, the variation in growth rates by number of observations is not large enough to make this re-weighting critically important.

The third column in the upper panel of Table VIII reports the average growth rate and estimated standard deviation estimated from the application of the jackknife to the data, i.e. estimating the model 26 separate times, each time removing one product from the sample. The mean jackknife point estimates and the jackknife estimate of their standard errors are incredibly close to those of the baseline lower left hand panel of Table VI earlier. With different relative price levels and trends, individual products will show unusually high or low levels and growth rates, but this distribution, with the adjustment of the 1st step covariance matrix, looks to be about what one expects from the normally distributed errors that underlie the specification of the baseline model. The delete-1 jackknifed growth rates range from .036 to .039 for the non-African economies and .030 to .036 for the African sample. This variation is smaller when growth is estimated using local income elasticities, as shown in the next section.

The fourth column in the top row of the upper panel of Table VIII provides an alternative calculation of means and standard errors using the bootstrap. My estimation procedure involves

multiple steps, with the calculations from earlier steps appearing as dependent variables or elements of the second step covariance matrices, while the survey data themselves are collected in clusters which are, typically, stratified by region, so the usual estimates of standard errors could be inaccurate.[32] Consequently, I bootstrap and recalculate all of the results 250 times, randomly sampling with replacement 135 surveys from my 135 surveys, randomly sampling the clusters within each survey (stratified by urban/rural location), and randomly sampling 26 from my 26 products. As shown in Table VIII, the resulting point estimates are close to those calculated using the original data, but the standard errors are between 30 and 60 % larger than those reported in the lower left hand panel of Table VI. The bootstrapped 95% confidence intervals for the non-African and African growth rates are .025 to .051 and .022 to .049, respectively. Given the enormous computational time involved, it is not possible to repeat this procedure for all of the other estimates I shall report, but this gives some sense of the degree to which the reported standard errors might be adjusted.[33]

Columns five and six of the upper panel of Table VIII re-estimate the first-step product demand equations using cluster random and fixed effects to explicitly allow for correlation in the error terms for households within clusters. When estimated with cluster random or fixed

[32]Given the complexities introduced by the sampling framework and the use of monte carlo estimates of the covariance matrix based upon the first step estimates, the standard two step formulas (e.g. Murphy and Topel 1985 or Hardin 2002) are not easily applied here. Outside of the bootstrap calculations, in all second step tables I report standard errors based upon the inverse of the negative hessian, while the first step covariance matrices (used in the monte carlo calculation of covariance matrices) use the sandwich adjustment for clustering.

[33]Lest there be any confusion, I should clarify that the difference between the third and fourth columns of Table VIII lies in the conceptualization of the sampling problem, and not in the jackknife vs. the bootstrap. The third column provides a non-parametric estimate of the variability induced by the sampling of products, given the first step estimates and the survey sample. The fourth column provides a non-parametric estimate of the variability induced by the sampling of surveys, clusters and products. I could just as easily bootstrap the third column, drawing 250 samples of 26 products from my 26 products. This produces the coefficients (s.e.) .034 (.005), .034 (.005), and .729 (.081). The jackknife, however, allows me to report the sensitivity of the growth rate to the extremes of the product growth distribution, as noted in the text.

effects[34], the first step quasi-income elasticities (i.e. coefficients on educational attainment) fall, implying that any movement in physical consumption levels is associated with greater real consumption growth. Consequently, the estimates of the growth and standard deviation of living standards are higher, as shown in the fifth and sixth columns of Table VIII. Although the cluster effects are always significant[35], 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, as reported in Table VI earlier.[36]

The lower panel of Table VIII explores the sensitivity of the results to alternative specifications of the probability model used in the estimation of the 1st step demands for the

[34]For the dichotomous variables, I use Butler & Moffitt's (1982) random effects specification, modeling 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. 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. The covariance matrix of the regional dummies and the estimated income elasticity are adjusted for the two-step procedure.

[35]The estimate of random cluster variation is always significantly different from zero, while a Hausman test of fixed versus random effects always concludes in favour of fixed effects, i.e. that there is correlation between the random effect and the independent variables.

[36]In all tables, when I don't have explicit cluster random or fixed effects I always adjust the first step covariance matrix (which is then used in the second step mle) for clustering.

dichotomous (0-1) variables. The plot of the probit (normal) cumulative density is, when rescaled, very similar to that of the logit, the weibull is asymmetric with somewhat fatter upper tails, the gompertz is asymmetric with fatter lower tails, the cauchy fattens both tails symmetrically, and the linear probability model produces thinner (zero) tails at the extremes of the distribution.[37] A fatter (thinner) tail means that changes in mean consumption levels in that region are associated with bigger (smaller) movements in the index determining the probability. Consequently, the gompertz and cauchy translate the observed movements in the low levels of sub-Saharan product consumption into higher estimates of aggregate consumption growth, while the weibull and linear model translate these movements into lower estimates of consumption growth. To resolve these differences, I apply the semiparametric discrete choice model developed by Gabler, Laisney and Lechner (1993), which uses a Hermite series expansion of the cumulative density, a flexible form that can approximate all of the other distributions used in the table.[38] As shown in the last column of the lower panel of Table VIII, this produces estimates that are just slightly below the baseline logit results of Table VI.

VI. Estimates Using Local Income Elasticities

The analysis above imposes the strong assumption that the return to education and the income response of demand is the same in all of the economies. Levels of development, however, are likely to affect both the return to education and the income elasticity of demand for particular products, while differences in local conditions and relative prices will not only

[37]Since the weibull and gompertz are asymmetric, the specification of a "success" (e.g. cough or no cough) affects the results. I adjust the measurement of the variables so that a success is associated with a positive quasi-income elasticity in the gompertz and a negative elasticity in the weibull. Thus, for example, success for the health variables is measured as no diarrhea, no fever, no cough, and child alive. Since the weibull and gompertz distributions are mirror images of each other, the opposite scaling simply exchanges the two sets of results.

[38]I set their k = 3, which results in the probability being the integral of a 6th order polynomial in XB times the normal density for XB.

influence levels of demand (as allowed above) but also income elasticities. Although I have explored a variety of functional forms, including semiparametric approximations, that translate a given coefficient on household educational attainment into different elasticities of demand at different levels,[39] it is still possible that heterogeneity across the sample accounts for my results. In particular, if the response of demand to educational attainment is systematically higher in sub-Saharan Africa or the return to education is systematically lower, then the estimates reported above will overstate African growth. In this section I address this concern by estimating demand patterns country by country and the return to education within and outside Africa. While I find heterogeneity across the sample, it is not systematically related to the results emphasized in this paper, i.e. with local demand coefficients I still find African growth to be the equal of non-African growth and close to four times as fast as reported in international sources.

(a) Methods

If one re-estimates the household demand equation (12) earlier above country by country, the resulting measures of regional living standards will be given by:

$$(14)' \quad \ln(\hat{C}^{R^c}_{prt}) = \hat{R}^c_E \left[\frac{\hat{a}_{prt}}{\hat{b}^c_p} + \hat{E}_{rt} \right]$$

where the superscript c on the quasi-income elasticity and the return to education emphasizes that these may now vary by country. These regional (i.e. urban/rural) measures can no longer be meaningfully compared across countries. However, the growth of product consumption within a country, translated into income equivalents with a constant country-specific income elasticity, can still be examined. Thus, I use population weights to produce country level measures

[39]Thus, for example, in the logit the elasticity of the purchase probability with respect to educational attainment is $(1-P)b_p$, where b_p is the product coefficient on educational attainment and P is the expected probability of purchase. Clearly, this falls as the consumption probability (level) rises.

$\ln(\hat{C}^{R^c}_{pct})$ (as in (15) earlier) and study the growth of these measures in the random effects regression:

$$(16)' \quad \ln(\hat{C}^{R^c}_{pct}) = a_{pc} + g_A t_A + g_{\sim A} t_{\sim A} + v_c t + v_p t + e_{pct} + \hat{e}_{pct}$$

Relative to equation (16) earlier, I now introduce a complete set of product x country dummies a_{pc} to account for the differing levels introduced by the country varying b^c_p, and make no attempt to compare overall country levels of consumption.

In the PWT a fixed set of international prices is used to weight local real expenditures, producing estimates of relative real consumption through space and time. In a similar fashion, the simplifying assumption of common international quasi-income elasticities of demand in the previous section allowed me to translate product consumption levels into income equivalents that could be compared internationally and intertemporally. In the national accounts, country specific constant price indices are used to calculate growth. From a welfare theoretic perspective, these produce more accurate measures of growth than the PWT (as the component real expenditures are weighted by the prices faced by the economic actors), but the resulting level measures are no longer comparable internationally. Similarly, in this section, in calculating the income equivalent of product consumption using local income elasticities, I produce measures of local growth that are theoretically (if not necessarily statistically) more accurate, but at the cost of no longer being able to compare levels internationally.

(b) First Step Estimates

As a preliminary, Table IX below runs separate Mincerian regressions for the African and non-African countries of log earnings from working for others on education and demographic

Table IX: Log Wage Regressions by Region						
		(1) survey dummies	(2) survey x rural/ urban dummies	(3) cluster random effects	(4) cluster fixed effects	(5) cluster fixed effects (IV)
Africa	educ	.140 (.003)	.129 (.003)	.123 (.002)	.113 (.003)	.139 (.009)
	age	.064 (.012)	.064 (.012)	.064 (.011)	.053 (.013)	.051 (.015)
	age2	-.001 (.000)	-.001 (.000)	-.001 (.000)	-.000 (.000)	-.000 (.000)
	sex	-.043 (.037)	-.056 (.037)	-.063 (.026)	-.061 (.030)	-.030 (.038)
	N	8041	8041	8041	8041	5897
~Africa	educ	.103 (.002)	.098 (.002)	.095 (.001)	.087 (.002)	.103 (.005)
	age	.042 (.008)	.042 (.008)	.046 (.007)	.051 (.008)	.050 (.010)
	age2	-.000 (.000)	-.000 (.000)	-.000 (.000)	-.001 (.000)	-.000 (.000)
	sex	-.548 (.019)	-.554 (.019)	-.553 (.019)	-.539 (.020)	-.555 (.023)
	N	14955	14955	14955	14955	12521
For notes and details on variable construction see Table IV and Appendix A. Coefficients on age2 are generally between -.0004 & -.0006 and significant.						

characteristics following the specifications described in Table IV earlier.[40] 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 estimated return to education of .139 in Africa and .103 outside of Africa, as the basis for my analysis.[41]

[40] As I do not have wage data for many countries, it is not possible to calculate a separate R_E for each country. The Africa/non-Africa breakdown employed above follows the results emphasized in the paper.

[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 IV), the woman's discount rises to 29% in Africa, while remaining at 59% for the non-African economies. However, the educational income profile, at .135 and .098 within and outside Africa, respectively, is largely unchanged.

Table X: 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	log(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
log Weight	.007 (.000)	.002 (.000)	51	Working (15-24)	-.009 (.007)	.044 (.005)	49
log Height	.002 (.000)	.001 (.000)	51	Working (25-49)	.052 (.010)	.067 (.007)	49
Diarrhea	-.035 (.004)	.023 (.003)	55	Birth (15-24)	-.014 (.003)	.014 (.003)	56
Fever	-.020 (.003)	.019 (.003)	55	Birth (25-49)	-.033 (.004)	.023 (.003)	56
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.

Table X describes the strong heterogeneity across countries in demand patterns. For each product I regress the first-step country level coefficients on household educational attainment on a constant. The figures reported in the table are the constant (mean country coefficient) and the standard error of the regression (standard deviation of the coefficients).[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 quite negatively associated with educational attainment in the Dominican Republic (coef. = -.10), where tap water is known to be contaminated.

[42]Since the dependent variables are estimated, I incorporate their covariance matrix in the likelihood. Thus, the constants are 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 XI: Growth Measures Based on Local Demand Patterns

$y_{pct} = a_{pc} + g_{\sim A}*t + g_A*t + v_p*t + v_c*t + e_{pct}$

first step logit for dichotomous variables

	All Products	Consumer Durables	Housing	Health	Family Economics
$g_{\sim A}$	.034 (.005)	.041 (.008)	.044 (.008)	.025 (.007)	.024 (.008)
g_A	.037 (.005)	.051 (.009)	.019 (.012)	.037 (.008)	.040 (.008)

	2^{nd} step w/out Σ(FS)	2^{nd} step w/out Σ(RS)	Jackknife Products	Bootstrap All Steps	1^{st} step cluster random effects	1^{st} step cluster fixed effects
$g_{\sim A}$	-.023 (.038)	.033 (.002)	.034 (.004)	.033 (.007)	.041 (.006)	.051 (.008)
g_A	.099 (.024)	.038 (.002)	.037 (.004)	.037 (.008)	.044 (.006)	.049 (.008)

alternative first step functional forms

	Probit	Weibull	Gompertz	Cauchy	Linear	Hermite
$g_{\sim A}$	.033 (.005)	.032 (.005)	.036 (.006)	.039 (.006)	.035 (.006)	.033 (.005)
g_A	.036 (.005)	.033 (.004)	.042 (.006)	.043 (.007)	.031 (.004)	.036 (.005)

Note: Unless otherwise noted, each specification is run separately for Africa and non-Africa, and includes product x country dummies (a_{pc}) and random effects for country and product growth (v_c, v_p). To save space I only report the estimated growth rates, g_i. . "w/out Σ(FS)": without first step estimation error covariance matrix in 2^{nd} step GLS covariance matrix. "w/out Σ(RS)": without covariance matrix induced by random shocks v_p and v_c in 2^{nd} step GLS covariance matrix.

(c) Second Step Growth Results

Table XI presents separate estimates of growth in the African and non-African economies based on equation (16)'. 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 somewhat higher than previously estimated in Table VI (.037 vs .034) and non-African growth somewhat lower (.034 vs .038). As before, the growth rates of durables are higher than the average, while African growth is substantially slower in housing. Growth in the non-traditional consumption measures, health and family economics,

is somewhat lower than the average, particularly outside of Africa, so these do not explain the discrepancy with international measures of growth.[43]

Turning to the results reported in the lower rows of the table, we see that estimates without adjustment for the precision of the first step estimates are nonsensical (methodologically and practically). The variation in the significance of first step estimates of the relationship between product consumption and education at the country level is enormous and accounting for this substantially re-weights the observations. In contrast, removing the adjustment for random variation in product and country growth rates has little effect on estimated growth. The estimated standard deviations of the product and country growth rates ($\sigma[v_p]$ and $\sigma[v_c]$) for the African (.017 and .012) and non-African (.018 and .013) economies in the upper-left hand cell of Table XI are substantial, but this re-weighting has little effect as there does not appear to be much systematic variation in growth rates by the number of observations within my unbalanced panel. A product jackknife produces means and standard errors that are close to those estimated under the baseline assumptions, showing once again that the covariance weighted product growth distribution approximates the normal distribution assumed in the baseline model. The gap between the slowest and fastest delete-1 jackknife growth rates is actually smaller than in previous section, i.e. ranging from .032 to .035 for the non-African countries and from .036 to .039 for sub-Saharan Africa. A bootstrap of all steps of the estimation process (surveys, clusters and products) suggests that the true standard errors might be about 40 to 60 % as large as those reported initially in the upper-left hand column of the table. The bootstrapped 95% confidence interval is .023 to .045 for non-African growth and .024 to .050 for sub-Saharan growth. As before, estimates with random and fixed effects yield higher average growth rates, alternative

[43]Removing these and focusing on durables and housing alone raises the non-African growth rate to .042 and the African growth rate to .038.

functional forms produce minor variation in the results, and a flexible hermite approximation returns growth estimates which are close to those of the baseline model.

All of these results follow the patterns reported in the previous section. There is, without a doubt, considerable heterogeneity across countries in demand patterns, but this averages out completely and does not eliminate the surprisingly high growth, particularly for sub-Saharan Africa, indicated by the DHS data.

VII. Conclusion

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 log consumption 98 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.

VIII. Bibliography

Becker, Gary S., Tomas. J. Philipson, and Rodrigo R. Soares. "The Quantity and Quality of Life and the Evolution of World Inequality." American Economic Review, Vol. 95, No. 1 (March 2005): 277-291.

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.

"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.

Gabler, Siegfried, Francois Laisney and Michael Lechner. "Seminonparametric Estimation of Binary-Choice Models with an Application to Labor-Force Participation." Journal of Business and Economic Statistics, Vol. 11, No. 1 (January 1993): 61-80.

Hardin, James W. "The Robust Variance Estimator for Two Stage Models." The Stata Journal (2002) Vol. 2, No. 3: 253-266.

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.

Johnson, Simon, William Larson, Chris Papageorgiou, and Arvind Subramanian. "Is Newer Better? Penn World Table Revisions and Their Impact on Growth Estimates." Working paper, October 2009.

Jones, Charles and Peter Klenow. "Beyond GDP? Welfare across Countries and Time." Manuscript, Stanford University, February 2011.

Murphy, Kevin M. and R.H. Topel. "Estimation and Inference in Two Step Econometric Models." Journal of Business and Economic Statistics (1985): 370-379.

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.

IX: Appendix A: 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 as follows:

> Constructed floor: hv213 <= 13 (dirt/sand/dung) = 0, otherwise (cement/wood/tiles/etc) = 1. Flush toilet: hv205 < 21 (including septic tanks) = 1, otherwise (pit/latrine/bush/etc) = 0. Tap drinking water: hv201 < 21 (tapped or piped) = 1, otherwise (well/stream/lake/etc) = 0. Diarrhea, fever and cough in past 2 weeks: yes answers 1 or 2 coded as 1 (extra detail on last 24 hours not universal across surveys and not used), no coded as 0. 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 V) are constructed as follows:

> Log 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.

For the wage regressions in Table IV, I restrict myself to female and male individuals aged 25 to 65 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).

Table AI: DHS and Associated Surveys Used in the Paper					
Benin	96*, 01, 06	Niger	92, 98, 06	Bolivia	94*, 98*, 03
Burkina Faso	92, 98, 03	Nigeria	90, 99*, 03	Brazil	91, 96
Cameroon	91, 98, 04	Rwanda	92, 00, 05	Colombia	90, 95*, 00, 05
Cen. Af. Rep.	94*	Senegal	92, 05	Dom. Rep.	91, 96*, 99, 02
Chad	96*, 04	South Africa	98*	Guatemala	95*, 98*
Comoros	96*	Tanzania	92, 96, 99, 03, 04	Guyana	05
Congo	05	Togo	98*	Haiti	94, 00, 05
Cote D'Ivoire	94, 98, 05	Uganda	95*, 00, 06	Honduras	05
Ethiopia	00, 05	Zambia	92, 96*, 01	Nicaragua	97*, 01
Gabon	00	Zimbabwe	94*, 99, 06	Paraguay	90
Ghana	93, 98*, 03			Peru	92, 96*, 00, 04
Guinea	99, 05	Bangladesh	93, 96, 99, 04		
Kenya	93, 98, 03	Cambodia	00, 05	Armenia	00, 05
Lesotho	04	India	92, 98, 05	Egypt	92, 95*, 00, 03, 05
Madagascar	92, 97*, 03	Indonesia	91, 94, 97, 02	Kazakhstan	95, 99
Malawi	92, 00, 04	Nepal	96*, 01, 06	Kyrgyz Rep.	97
Mali	95*, 01, 06	Pakistan	90	Moldova	05
Mozambique	97*, 03	Philippines	93, 98*, 03	Morocco	92, 03
Namibia	92, 00	Vietnam	97, 02	Turkey	93, 98*, 03
				Uzbekistan	96

Notes: Years denote date when survey began; data collection often continues into the following year. 9x numbers are 199x, 0x numbers are 200x. (*) Surveys with wage income data

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), eliminating most of the 15-24 age group. 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. 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.

X: Appendix B: Random Variation and Observation Weights

In equation (16) I allow for random variation in the level of consumption at the product x country level (u_{pc}) and the trends of particular products or countries (v_p, v_c). In this appendix I explain the claim in the text that these random shocks affect the weighting of observations in the estimation of the product and country fixed effects a_p and a_c and the time trends g_A and $g_{\sim A}$.

I begin by describing the solution to a standard problem. Consider the panel regression:

$$\text{(B1)}\quad Y_{it} = X_{it}\beta + u_i Z_{it} + \varepsilon_{it}$$

where i denotes the panel data group and t the within group observation, u_i is a group specific shock multiplied by the variable Z_{it}, and X_{it}, β and Z_{it} are 1xk, kx1 and 1x1, respectively. The covariance matrix for the T_i observations for group i is given by:

$$\text{(B2)}\quad \Sigma_i = \sigma_\varepsilon^2 I_{Ti} + \sigma_u^2 Z_i Z_i^{/}$$

where I_{Ti} is the identity matrix of dimension T_i, and Z_i is the column vector of Z_{it} observations for group i. Letting X_i and Y_i denote the corresponding matrices of T_i observations for X_{it} and Y_{it}, following standard GLS results the maximum likelihood estimate of β is given by:

$$\text{(B3)}\quad \hat{\beta} = \left(\sum_i X_i^{/}\Sigma_i^{-1}X_i\right)^{-1}\left(\sum_i X_i^{/}\Sigma_i^{-1}Y_i\right) = \left(\sum_i \tilde{X}_i^{/}\tilde{X}_i\right)^{-1}\left(\sum_i \tilde{X}_i^{/}\tilde{Y}_i\right),$$

$$\text{where}\quad \tilde{X}_i = \Omega_i^{-1/2}X_i,\ \ \tilde{Y}_i = \Omega_i^{-1/2}Y_i,\ \ \Omega_i^{-1/2}\Omega_i^{-1/2} = \Sigma_i^{-1}\quad \text{and}$$

$$\Omega_i^{-1/2} = \frac{1}{\sigma_\varepsilon}\left[I_{Ti} - \frac{\theta_i Z_i Z_i^{/}}{Z_i^{/}Z_i}\right]\ \text{with}\ \ \theta_i = 1 - \frac{\sigma_\varepsilon}{\left(\sigma_\varepsilon^2 + \sigma_u^2 Z_i^{/}Z_i\right)^{1/2}}.$$

The dependent variable in (16) is indexed by three characteristics (product x country x time) and there are multiple random shocks on the right hand side. Some intuition into how the random shocks relate to the estimation of different coefficients can by arrived at by linking the three characteristics to the standard i x t notation and considering segments of the problem in isolation. With regards to the random effects u_{pc}, let i denote the product x country grouping and t denote the time dimension, with Z_{it} equal to the constant 1. Further, considering only the estimation of either the country or product fixed effects (a_c or a_p), let X be the k mutually exclusive 0/1 indicator variables for the product or country categories. Since the X variables are orthogonal to each other, the cross-product matrices are diagonal and, applying (B3), we see that the estimate of the coefficient for the kth group is given by:

$$\text{(B4)}\quad \hat{\beta}_k = \left(\sum_{i\in S(k)}(1-\theta_i)^2 T_i\right)^{-1}\left(\sum_{i\in S(k)}(1-\theta_i)^2\sum_{t\in S(i)}Y_{it}\right)\ \text{with}\ \ \theta_i = 1 - \frac{\sigma_\varepsilon}{\left(\sigma_\varepsilon^2 + \sigma_u^2 T_i\right)^{1/2}},$$

where *S(k)* is the set of i (product x country) groupings appearing in category k and *S(i)* is the set of t (time) observations for grouping i. The OLS estimate of the kth fixed effect equals (B4) with θ_i equal to 0 for all i. As θ_i is larger for groups with a larger number of observations T_i, we see

that relative to OLS the GLS estimate places less than a one-for-one weight on observations from larger groups. This explains my claim regarding the influence of u_{pc} on the country or product fixed effects in equation (16) (a_c and a_p).

Regarding the random variation in trends, v_p and v_c in (16), let i be the country or product (respectively), t the cross of the remaining categories (i.e. product x time or country x time), and Z_{it} and X_{it} the year of the observation (say, yr_{it}). Thus, in this case I am considering (B1) as a univariate regression on a time trend (without a constant) with random variation across groups in the trend. Applying (B3) we find that

$$\text{(B5)}\quad \hat{\beta}_{yr} = \left(\sum_i (1-\theta_i)^2 \sum_{t \in S(i)} yr_{it}^2 \right)^{-1} \left(\sum_i (1-\theta_i)^2 \sum_{t \in S(i)} Y_{it} yr_{it} \right) \text{ with } \theta_i = 1 - \frac{\sigma_\varepsilon}{\left(\sigma_\varepsilon^2 + \sigma_u^2 \sum_{t \in S(i)} yr_{it}^2 \right)^{1/2}}.$$

The OLS estimate of the time trend equals (B5) with θ_i equal to 0 for all i. If the magnitude of the yr observations is roughly the same across i groups, the sum of their squares will be roughly proportional to T_i, so θ_i will be larger for groups with more observations. Once again, we see that relative to OLS the GLS estimate places less than a one-for-one weight on observations from larger groups. This explains my claim regarding the influence of v_p and v_c on the estimation of the time trends g_A and $g_{\sim A}$ in equation (16).

I introduce the random variation u_{pc} to allow for permanent differences in consumption levels brought about by relative price differences, and the variation v_p and v_c to allow for the fact that different products (because of global price trends) and countries have different trend growth rates. In the actual estimation of (16) all of the random shocks and coefficients are estimated simultaneously, which introduces interactions not explored in the equations above, but I believe these examples provide some intuition as to how these effects influence the coefficient estimates.