NBER WORKING PAPER SERIES

LIGHTS, CAMERA,... INCOME!: ESTIMATING POVERTY USING NATIONAL ACCOUNTS, SURVEY MEANS, AND LIGHTS

Maxim Pinkovskiy
Xavier Sala-i-Martin

Working Paper 19831
http://www.nber.org/papers/w19831

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2014

We would like to thank Richard Crump for useful suggestions. Any views expressed in this paper are the authors and do not necessarily reflect those of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research. All errors are our own.

Lights, Camera,... Income!: Estimating Poverty Using National Accounts, Survey Means, and Lights
Maxim Pinkovskiy and Xavier Sala-i-Martin
NBER Working Paper No. 19831
January 2014
JEL No. D31,E01,O1,O4

ABSTRACT

In this paper we try to understand whether national accounts GDP per capita or survey mean income or consumption better proxy for true income per capita. We propose a data-driven method to assess the relative quality of GDP per capita versus survey means by comparing the evolution of each series to the evolution of satellite-recorded nighttime lights. Our main assumption, which is robust to a variety of specification checks, is that the measurement error in nighttime lights is unrelated to the measurement errors in either national accounts or survey means. We obtain estimates of weights on national accounts and survey means in an optimal proxy for true income; these weights are very large for national accounts and very modest for survey means. We conclusively reject the null hypothesis that the optimal weight on surveys is greater than the optimal weight on national accounts, and we generally fail to reject the null hypothesis that the optimal weight on surveys is zero. Using the estimated optimal weights, we compute estimates of true income per capita and $1/day poverty rates for the developing world and its regions. We get poverty estimates that are substantially lower and fall substantially faster than those of Chen and Ravallion (2010) or of the survey-based poverty literature more generally.

Maxim Pinkovskiy
Federal Reserve Bank of New York
33 Liberty Street
New York, NY 10045
maxim.pinkovskiy@ny.frb.org

Xavier Sala-i-Martin
Department of Economics
Columbia University
420 West 118th Street, 1005
New York, NY 10027
and NBER
xs23@columbia.edu

1 Introduction

Our understanding of how rapidly poverty is falling in the developing world critically depends on how we measure the means of country income distributions. Disagreement over whether this variable is best captured by GDP per capita from the national accounts or by mean consumption from household surveys forms the crux of the differences between researchers asserting that world poverty has fallen dramatically and has ceased to be a major presence in the developing world outside of Africa, and researchers suggesting that it has declined more modestly, and remains a problem to be grappled with. Thus, Bhalla (2002), Sala-i-Martin (2002, 2004, 2006), Pinkovskiy and Sala-i-Martin (2009) and Sala-i-Martin and Pinkovskiy (2010) use national accounts data to find that world poverty has declined to 13% of the developing world population by 2000 (Bhalla 2002) or to less than 6% of the developing world population by 2006 (Pinkovskiy and Sala-i-Martin 2009), and that Africa is on track to halve its 1990 level of poverty within a few years of 2015. On the other hand, Chen and Ravallion (2001, 2004, 2010) find that world poverty was 25% in 2005 (down from 52% in 1992), that the number of the poor (though not the fraction) continues to increase, and that the developing world outside China (in particular, Africa) is not on track to achieve the Millennium Development Goals. The major difference in methodology between these two sets of studies is estimation of the means of country income distributions: survey means have a much lower level and a slower growth rate than do national accounts-based GDP estimates, and this difference dwarfs any difference in poverty estimates that can be attributed to differing parametric or nonparametric assumptions about the course of within-country income inequality. Deaton (2005) discusses the sources of this discrepancy, some working to bias national accounts and others to bias survey means, and Young (2012) argues that national accounts (and, a fortiriori, survey means) underestimate economic growth in Africa based on consumption data from the Demographic and Health Surveys, but so far, to our knowledge, there has been no success in reconciling national accounts and survey means and in showing which source of data is superior.

In this paper, we hope to contribute to the literature by proposing a way to assess whether national accounts or survey means perform better in capturing differences in income across countries and over time, creating a new measure of income per capita that is an optimal combination of national accounts and survey means data, and presenting estimates of world poverty from 1992 to 2010 using this measure. Our main idea is to exploit a third, independently collected source of data on economic activity around the world: satellite-recorded nighttime lights (Elvidge et al. 1997). It is intuitive that nighttime lights should reflect economic activity to some degree because light is a critical input in many production processes and consumption activities (e.g. outdoor lighting, consumption activities at night in private homes or public places, transportation of goods and people, productive activity in factories and office buildings, and evening consumption of mass media). The main advantage of using nighttime lights rather than a different proxy for income is that the data generating process for lights allows us to distinguish the components of national accounts (or survey means) that reflect true income rather than measurement error. In general, a positive correlation between measured income (national accounts or survey means) and nighttime lights could be due to two factors: that they are both correlated with true income, or that their measurement errors are strongly correlated with each other. However, the latter possibility is implausible because the generating process of nighttime lights data is to a very large degree independent of the generating process either of national accounts or of survey means. For example, measured income is collected by statisticians interacting with survey respondents, while nighttime lights are recorded impersonally by satellites. Statistical teams use different procedures in different countries, while lights are recorded homogeneously across national borders. Both national accounts and survey means may suffer from nonrandom nonresponse and misreporting, whereas nighttime lights do not require compliance or truthfulness of the surveyed population to record whatever lights exist. Moreover, nighttime lights may vary because of climatic conditions such as auroral activity, cloudiness and humidity, or because of cultural attitudes towards lighting, which presumably do not affect measurement errors in national accounts or survey means. Therefore, the strength of the correlation between nighttime lights and measured income is directly related to the strength of the correlation between the given income measurement and the true income it is trying to measure. We can use the ratios of correlations between nighttime lights and different income measurements to assess the relative strengths of the correlations between these income measurements and unobserved true income.

Our goal in this paper is twofold: first, test whether national accounts or survey means better reflect variation in true income across countries and over time, and second, create a new proxy for true income that will allow us to compute poverty rates in developing countries and assess the evolution of world poverty. We find that under our assumptions, the national accounts GDP data reflect variation in income per capita much better than survey means do. If we wish to construct an optimal loglinear combination of national accounts and survey means as an improved proxy for true income per capita, we find that the weight that we wish to place on survey means is 18% of the weight that we wish to place on national accounts GDP. This is very different from prior methods of combining survey means and national accounts, which have used Bayesian theory and the principle of insufficient reason to assign equal weights to survey means and their predicted value based on national accounts GDP; hence survey means got more than 100% of the weight placed on national accounts (Chen and Ravallion, 2010). This conclusion also does not change whether we look at predicting cross-country differences or growth rates of true income, or when we allow for the relationships between the measures of income we consider (GDP, survey means and lights) to be affected by other variables, or to vary across space and over time.

We can use this methodology to compute optimal loglinear predictors of true income in terms of national accounts and survey means and estimate poverty for the world as a whole over time. The precise magnitude of our poverty estimates depends on assuming a scale for the unobserved true income measure. Under the plausible assumption that this scale is at its long-run value given the weights that should be placed on national accounts and survey means, we find that poverty in the developing world is very close in level and in trend to the national accounts-based measurements. Even if we use the scaling assumption that is most favorable for replicating poverty estimates obtained with survey means (Chen and Ravallion 2001, 2004, 2010) we find that poverty is lower and has declined by more than has been found by research using survey means alone, the difference being statistically significant if we account for the statistical error in our computation of the optimal weights. This result is also robust to flexible specifications of the relationships between the different measures of income, or to accounting for the potential mismeasurement (and specifically, underestimation) of the growth in inequality, rather than just growth in mean income, in the surveys.

In this paper, we remain agnostic about the precise reasons for which national accounts appear to be a superior measure of true income than survey means are. However, an explanation that is consistent with the literature exploring the discrepancies between national accounts and survey means (Bhalla (2002), Deaton (2005, 2010)) is that national accounts-based measures of GDP better capture income streams arising from consumption of owner-occupied housing and from public goods provided by the government than household surveys do. This explanation is particularly consistent with our results because nighttime lights would be expected to capture these two sources of income particularly well, as housing stock and public goods are major sources of private and public lighting. It is also unlikely that nighttime lights capture aspects of public or private consumption that do not figure substantially in the consumption of the poor, such as national defense spending, spending on luxury goods, transfer of money abroad and corruption.

The rest of the paper is organized as follows. Section 2 presents a review of the relevant literatures. Section 3 describes the lights measure that we use. Section 4 describes our mathematical framework for computing optimal weights and states the assumptions that we make on the data generating processes for lights, GDP and surveys. Section 5 presents our results for relative weights. Section 6 presents our poverty estimates for the world and for some of its regions. Section 7 concludes.

2 Related Literature

Our paper is most directly in the literature on estimating the world distribution of income, which is divided into two major strands: papers estimating means of income distributions with national accounts

GDP [Bourguignon and Morrisson (2002), Bhalla (2002), Sala-i-Martin (2002, 2004, 2006), Pinkovskiy and Sala-i-Martin (2009), Sala-i-Martin and Pinkovskiy (2010), Pinkovskiy (2013)], and papers estimating means of income distributions with income or consumption survey means [Chen and Ravallion (2001, 2004, 2010), Milanovic (2005)]. All papers use data from household surveys to estimate within-country inequality in the distribution of income or consumption, which is also the procedure we will follow here. It is apparent (Pinkovskiy and Sala-i-Martin 2009, Dhongde and Minoiu 2010) that the most important factor explaining the differences between various world poverty estimates is whether national accounts GDP or survey means are used to anchor country income distributions, rather than the way in which the shape of these distributions is extracted from limited inequality data. Alvaredo and Gasparini (2013) provide a general review of this literature.

We are not the first to combine national accounts GDP and survey means in an attempt to obtain a better measure of true income. Deaton (2001) notes that if we have two faulty measures of the same thing, it is reasonable to try to use a weighted average of them to reduce their flaws. Chen and Ravallion (2010) draw on Karshenas (2003) to present a mixed method in which log survey means are averaged with their predicted value based on national accounts consumption. However, using a simple average of these two measures is arbitrary when it is possible to exploit auxiliary information to create an optimal, data-driven combination of national accounts and survey means, which is our contribution in this paper.

Many arguments have been made about the virtues and defects of national accounts and survey means. On the one hand, it is obvious that surveys suffer from nonresponse bias, which may have been growing over time (Bhalla (2002)). It is also the case that surveys may measure certain categories of spending, which may have been growing in importance as a share of consumption, incorrectly, such as spending on new goods (Bhalla 2002) or spending on public goods. On the other hand, it is plausible that household surveys, which are typically carried out by the World Bank itself, may be better implemented than the national accounts collection in developing countries. National accounts estimates are often constructed under assumptions that are implausible for many markets in developing countries (e.g. perfect competition), which may lead to overstating income through the inclusion of rents as value added (Deaton 2005). Moreover, survey nonresponse is unlikely to be independent of respondent income, with rich people in developing countries probably unlikely to respond to surveys, or to reveal their incomes. For example, Korinek et al. (2005) finds that rich people in America are nearly 50% less likely to respond to surveys as poor people are (but Bhalla (2002) finds that consumption of luxuries is not substantially more underreported in India's 1993-1994 National Statistical Survey than is consumption of necessities). While it is not theoretically necessary that increasing nonresponse with income should decrease measured inequality (Deaton (2005) exhibits an admittedly special model in which nonresponse by the rich leaves inequality unchanged and decreases the survey mean only), there is the possibility that nonrandom nonresponse, growing over time, may mask rising inequality in developing countries.[1]

We believe that our analysis can avoid many of the pitfalls of either national accounts or survey means. Given that light is such an essential input to most meaningful economic activities, it is unlikely that our lights measure can be critiqued for attributing spurious or deleterious activites, such as monopoly rent extraction, to economic growth. Nor is it plausible to believe that the part of income that varies with light intensity is particularly unequally distributed, since light intensity derives from agglomeration of multiple lit structures, which are unlikely to be very closely owned. We think that nighttime lights most likely reflect lighting in houses, production facilities (stores, factories, ports) and modes of transportation. Since nighttime lights data is collected through an impersonal, nonintrusive process, concerns about nonresponse do not apply. While we cannot rule out theoretically that surveys underestimate inequality as well as economic growth, in our analysis, we can perform robustness checks by assuming counterfactual paths for the growth rate of the Gini coefficient over time; we find only modest effects on poverty estimates for very large mismeasurements of the Gini coefficients.

More recently, several papers have challenged the quality of national accounts, especially for the estimation of growth rates. Johnson et al. (2009) find large discrepancies between different vintages of the Penn World Tables, and show that successive updates of the PWT may yield different answers to questions on year-to-year variation in country growth rates. Young (2012) finds that changes in consumption of common goods recorded in the Demographic and Health Surveys in a variety of developing countries imply much higher growth rates in income than are recorded in the national accounts. Our paper also tries to assess the quality of and improve national accounts data by using external information (as in Young [2012]). We sidestep the difficulties of using the PWT by using PPP-adjusted GDP per capita from the World Development Indicators, but our focus is ultimately on comparing national accounts and survey means to each other rather than on assessing both of them against a third measure.

There is a large and growing literature that uses satellite data on lights at night as a proxy for income. Elvidge et al. (1997, 1999, 2007, 2010, 2012) describe the nighttime lights data, show that nighttime lights correlate well with electricity utilization and GDP, and construct poverty and inequality measures based on nighttime lights. Henderson, Storeygard and Weil (2009, 2012) (the latter, hereafter, HSW (2012)) and Chen and Nordhaus (2010) (hereafter CN (2010)) argue that nighttime lights are systematically correlated with rates of economic growth and may improve measures of GDP for poor countries. Michalopoulos

[1] Survey estimates of disposable income from the Luxembourg Income Study (LIS) (LIS 2013) find mean incomes to be larger and Gini coefficients to be smaller for the several developing countries and years for which both LIS estimates and survey estimates used in Chen and Ravallion (2010) are available. For example, the LIS survey for Brazil finds that mean disposable income is \$6000 and the Gini is 48; the Brazilian survey cited by Chen and Ravallion (2010) finds that mean income is \$3900 and the Gini is 56. Comparisons for Colombia, Estonia, Guatemala, Hungary, Mexico, Peru and Poland are similar.

and Papaioannou (2012, 2013), Alesina, Michalopoulos and Papaioannou (2013) and Pinkovskiy (2013) use nighttime lights to construct measures of income within African ethnicities, world ethnicities and neighborhoods of national borders respectively. Our paper is closest in spirit to HSW (2012) and CN (2010) in that it also considers the problem of optimally combining measures of economic activity; however, instead of using nighttime lights as a component of such a measure, we use it as an auxiliary variable to help uncover the correlation structure between the measures we do wish to use in our index. We also consider a different type of predictor for true income that do either HSW (2012) or CN (2010), which allows us to make fewer assumptions on the data generating processes that we consider.

3 The Nighttime Lights Measure

Data on luminosity at night is collected by the DMSP-OLS satellite program and is maintained and processed by the National Oceanic and Atmospheric Administration (NOAA). Satellites orbit the Earth every day between 20:30 and 22:00, sending images of every location between 65 degrees south latitude and 65 degrees north latitude at a resolution of 30 arcseconds (approximately 1 square km at the equator). The images are processed to remove cloud cover, snow and ephemeral lights (such as forest fires) to produce the final product available for download at

http://www.ngdc.noaa.gov/dmsp/downloadV4composites.html

The nighttime lights data is available from 1992 to 2012, and we use the data up to 2010 because of the paucity of household surveys after that date that have already been made available for research.

Each pixel (1 square kilometer) in the luminosity data is assigned a digital number (DN) representing its luminosity. The DNs are integers ranging from 0 to 63, with the relationship between DN and luminosity being

$$\text{Radiance} \propto \text{DN}^{3/2}$$

(Chen and Nordhaus 2010). However, pixels with DN equal to 0 or 63 may be top- or bottom-censored. Another known problem with the lights data is the presence of overglow and blooming: light tends to travel to pixels outside of those in which it originates, and light tends to be magnified over certain terrain types such as water and snow cover (Doll 2008). Given that we will compute national-level estimates of aggregate lights, it is unlikely that these sources of error will be large enough or sufficiently correlated with important variables that they will confound our analysis. Another problem may be that satellites age in space and are eventually retired. Hence, they might give inconsistent readings from year to year, or new satellites may give fundamentally different readings from old ones.[2] While some evdence of this problem exists, we will show in Sections 5 and 6 that our estimates of the optimal ways of combining national accounts and survey means are almost invariant to allowing the relationship between national accounts, survey means and lights to differ from year to year.

In our analysis, we will use the nighttime lights to construct an aggregate radiance measure for each country in each year and use it as a proxy for aggregate income. We construct this measure by computing the radiance within each pixel in each country and adding up the pixels. Using alternative aggregation formulas (for instance, adding up the DN's across pixels) yields very similar results. For years with multiple satellites available, we average the logarithms of our aggregate luminosity measure, following HSW (2012).

It is well established that lights are very well correlated with national accounts GDP, in levels, growth rates and business cycle fluctuations. Henderson, Storeygard and Weil (2012) provide these correlations, as well as dramatic pictures of long-term differences in incomes (North vs. South Korea) as well as short-term fluctuations (the Asian financial crisis of 1997-8) reflected in lights. We provide two pictures emphasizing two poor countries for which national accounts and survey means give completely different growth estimates: Angola and India. Figure I presents a picture of nighttime lights over southern Africa in 2000 and in 2009. We see that Angola, which according to the household surveys has experienced a 5% *decline* in per capita income, and according to the national accounts has experienced a *doubling* of per capita income (for all statistics, see Appendix Table AII), has many more lights in 2009 than it did in 2000. Most other southern African countries also have more lights in 2009 than in 2000 (Botswana, Zambia, Mozambique, South Africa, Malawi), but Zimbabwe has fewer lights, because of its economic collapse under the disastrous hyperinflationary policies of Robert Mugabe. Figure II gives a view of India between 1994 and 2010. According to household surveys, its per capita income grew by 29% over this time period, but according to the national accounts its per capita income more than doubled. We see that lights in India increase dramatically both in their intensity over the major cities as well as in their extent over previously unlit areas of the country (for example, there is a marked increase over Bihar in the Ganges valley, one of India's poorest provinces). The increase in lights in India and in Angola is much closer to what is suggested by the national accounts than by the survey means, and it is very unlikely that this additional economic activity benefits exclusively the rich because of the spatial extent of the new lights. Moreover, the fact that we observe Angola growing and Zimbabwe shrinking over the same period of time and measured with the same satellites suggests that uniform differences in

[2] The satellites from which data is available are: F10 (1992-1994), F12 (1994-1999), F14 (1997-2003), F15 (2000-2007), F16 (2004-2009) and F18 (2010-).

brightness between satellites do not dominate the observed variation in lights. While these figures are only suggestive (the lights we observe are aggregate rather than per capita lights), they already provide a hint that economic growth in the developing world may have been more extensive than surveys show.

4 Mathematical Framework

4.1 Calculation of Relative Weights in Optimal Forecasts

Consider the following model of our data. We have $N+1$ candidate proxies y_i^n, $n = 0, ..., N$ for log true income, denoted y_i^*. We also have a vector of covariates x_i of length K (which always includes a constant but may also include other variables). Define the loglinear forecast of y_i^* as

$$z_i = \eta\left(X_i\right) + \gamma' y_i$$

where y_i is a vector of the y_i^n's, X_i is an $N \times K$ matrix of the x_i's, η is a linear function, and γ is a vector of weights.

To fix notation, we set the log lights-based GDP measure to be y_i^0, log World Bank GDP per capita to be y_i^1, log survey means to be y_i^2 and other GDP-based measures (if any) are y_i^3, y_i^4 etc.

We are interested in two quantities. First, we wish to assess the weight given to log survey means (y_i^2) in the optimal forecast relative to the weight given to log World Bank GDP per capita (y_i^1). This is given by

$$\hat{\omega} := \hat{\gamma}^2 / \hat{\gamma}^1$$

where $\hat{\gamma}$ is the optimal weight vector.

We are also interested in computing values for z_i itself for all countries and years in our sample and in using z_i in place of y_i^1 or y_i^2 as the logarithm of the true mean of the income distribution for the country and year corresponding to observation i. Doing this will require more assumptions than calculating $\hat{\omega}$, but our conclusions will be qualitatively robust to a variety of alternatives for the assumptions we have to add.

To calculate $\hat{\omega}$ we make the following assumptions:

$$y_i^n = \alpha_n\left(x_i\right) + \beta_n y_i^* + \varepsilon_i^n,\ \varepsilon_i^n \text{ i.i.d. across } i. \tag{A1}$$

$$E\left(\varepsilon_i^n \varepsilon_i^m | X_i\right) = \sigma_{nm} \tag{A2}$$

$$E\left(\varepsilon_i^n y_i^* | X_i\right) = 0 \tag{A3}$$

$$E\left(\varepsilon_i^n \varepsilon_i^0 | X_i\right) = 0 \tag{A4}$$

All of these assumptions have been made (without conditioning on controls) in the previous literature, notably by Henderson, Storeygard and Weil (2012) and Chen and Nordhaus (2010). The iid assumption is the substance in Assumption A1. Assumption A2 assumes homoskedasticity with respect to X_i. Assumption A3 mandates that the error in each proxy is an affine function of true income plus noise that is uncorrelated with income, and that the linear relationship is stable across the sample. Assumption A4 is the key reason for the use of the lights data: it says that the random errors in lights measurement are uncorrelated with the random errors in GDP or survey-based income measurement. This assumption has also been made in HSW (2012) and CN (2010). This is a plausible assumption because the data generating processes of the lights data and of GDP (or surveys) are largely disjoint; lights data is collected by satellites without respect for borders, institutional structures, or people's desire to respond to surveys, whereas GDP and survey data are obtained primarily by asking people, who may be unwilling to respond accurately, if at all. There is a fear that errors in GDP, surveys and lights have a common component; for instance, if the product of different industries is differentially accounted for by surveys (because of differential ease of reporting), and also generates different amounts of light per unit of true income, then the errors in both light and GDP will have a common component (though likely with different coefficients, or even signs). However, since β_n is not necessarily unity, all our GDP proxies are allowed to have a bias that is affine in true income, so if differential industrial composition causes a bias that is related to GDP size (which is not an implausible assumption, at least to first order) this will be reflected in the β_n's not being equal to unity. Moreover, the framework accomodates including controls that may help make this assumption more credible.

Suppose that the parameters $\alpha = [\alpha_n]_{n=0}^N$, $\beta = [\beta_n]_{n=0}^N$, and $\Sigma = [\sigma_{nm}]_{n=0,m=0}^{N,M}$ are known. Then, the difference between the proxy z_i and y_i^* can be expressed as follows:

$$\begin{aligned} z_i - y_i^* &= \eta\left(X_i\right) + \gamma' y_i - y_i^* \\ &= \eta\left(X_i\right) + \gamma' \alpha\left(X_i\right) + \left(\gamma'\beta - 1\right) y_i^* + \gamma' \varepsilon_i \end{aligned}$$

Note that if we set

$$\begin{aligned} \eta\left(X_i\right) &= -\gamma'\alpha\left(X_i\right) \\ \gamma'\beta &= 1 \end{aligned} \quad \text{(C)}$$

then our proxy z_i will be unbiased for all values of X_i, regardless of the functional form of $E\left(y_i^*|X_i\right)$. The mean squared error of z_i under Assumptions A1 and A2 is given by

$$\begin{aligned} E\left(\left(z_i - y_i^*\right)^2 |X_i\right) &= \left(\eta\left(X_i\right) + \gamma'\alpha\left(X_i\right) + \left(\gamma'\beta - 1\right) E\left(y_i^*|X_i\right)\right)^2 \\ &\quad + \left(\gamma'\beta - 1\right)^2 \sigma_*^2 + \gamma'\Sigma\gamma \end{aligned} \quad \text{(MSE)}$$

Consider the γ that minimizes (MSE) subject to the unbiasedness constraint (C). This γ solves the simplified program

$$\hat{\gamma} = \arg\min_{\gamma} \gamma'\Sigma\gamma \text{ subject to } \gamma'\beta = 1 \ \ (\lambda)$$

since $\eta\left(X_i\right)$ imposes no restrictions on γ. By taking first order conditions, we get the system of equations

$$\begin{pmatrix} \Sigma & \beta \\ \beta' & 0 \end{pmatrix} \begin{pmatrix} \hat{\gamma} \\ \lambda \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$$

which imply that

$$\hat{\gamma} = \left(\beta'\Sigma^{-1}\beta\right)^{-1} \Sigma^{-1}\beta$$

If we relax the unbiasedness constraint, and set

$$\eta\left(X_i\right) = -\left(\gamma'\alpha\left(X_i\right) + \left(\gamma'\beta - 1\right) E\left(y_i^*|X_i\right)\right)$$

the optimal solution solves

$$\tilde{\gamma} = \min_{\gamma} \left(\gamma'\Sigma\gamma + \left(\gamma'\beta - 1\right)^2 \sigma_*^2\right)$$

and is given by

$$\tilde{\gamma} = \left(\Sigma + \beta\beta'\sigma_*^2\right)^{-1}\beta\sigma_*^2$$

Under assumptions A1-A4 we cannot solve for the optimal weight vectors $\hat{\gamma}$ and $\tilde{\gamma}$, but we can compute the ratios $\frac{\hat{\gamma}_n}{\hat{\gamma}_m}$ and $\frac{\tilde{\gamma}_n}{\tilde{\gamma}_m}$ (which turn out to be the same) for any $n, m \neq 0$ (that is, for the relative weights of any two proxies excluding the lights proxy). We can define the variance-covariance matrix S of the residuals

$$\tilde{y}_i^n = E\left(y_i^n|X_i\right)$$

and note that

$$\begin{aligned}
S_{n,n} &: = var\left(y_i^n|X_i\right) = \beta_n^2\sigma_*^2 + \sigma_n^2 \\
S_{n,0} &: = cov\left(y_i^n, y_i^0|X_i\right) = \beta_n\beta_0\sigma_*^2 \\
S_{n,m} &: = cov\left(y_i^n, y_i^m|X_i\right) = \beta_n\beta_m\sigma_*^2 + \sigma_{nm} \text{ for } n, m \geq 1
\end{aligned}$$

where the left hand-sides are known data elements (they are entries of the variance-covariance matrix S) and the right hand-sides are equations in β, Σ and σ_*^2.

Note that the equations for S_{n0} use Assumption A3 and are key towards identifying ratios of the parameters $\beta_0, \beta_1, ..., \beta_n$. They are the algebraic statement of the inference we draw from assuming that the measurement error in lights is uncorrelated with the measurement error in the measured income proxies: any covariance between lights and measured income is proportional to β_n, the proportionality constant being $\beta_0\sigma_*^2$. If we consider a positive covariance between national accounts and survey means (y_i^1 and y_i^2) then we cannot reach the same conclusion: $cov\left(y_i^n, y_i^m|X_i\right)$ may be large because $\beta_n\beta_m$ is large or because σ_{nm} is large. Since Assumption A3 rules out a σ_{n0} term, it allows us to estimate the ratio β_n/β_m for any n and $m \geq 1$, and thus to identify the relevant parameters in our model.

Then,

$$\begin{aligned}
\frac{\beta_0}{\beta_1} &= \frac{S_{1,0}}{\beta_1^2\sigma_*^2} \\
\frac{\beta_n}{\beta_1} &= \frac{S_{n,0}}{S_{1,0}} \\
\sigma_n^2 &= S_{n,n} - \beta_n^2\sigma_*^2 \\
\sigma_{nm} &= S_{n,m} - \beta_n\beta_m\sigma_*^2
\end{aligned}$$

More compactly, we can write

$$\hat{\beta}\left(\sigma_*^2\right) := \frac{1}{\beta_1}\beta = \left[\begin{array}{c} \frac{S_{1,0}}{\beta_1^2\sigma_*^2} \\ \frac{\hat{C}}{S_{1,0}} \end{array}\right]$$

where $\hat{C} = [S_{1,0}, S_{2,0}..., S_{n,0}]'$. Hence, all the coefficient ratios β_n/β_m for any n and $m \geq 1$ are identified. We can also note that under Assumption A3

$$S = \Sigma + \beta\beta'\sigma_*^2$$

Noting that by the binomial inverse theorem,

$$\begin{aligned}
\Sigma^{-1} &= \left(S - \beta\beta'\sigma_*^2\right)^{-1} = S^{-1} - \frac{\sigma_*^2}{1+\beta' S^{-1}\beta\sigma_*^2} S^{-1}\beta\beta' S^{-1} \\
&\Rightarrow \beta'\Sigma^{-1}\beta = \frac{\beta' S^{-1}\beta}{1+\beta' S^{-1}\beta\sigma_*^2} \\
&\Rightarrow \Sigma^{-1}\beta = S^{-1}\beta\left(\frac{1}{1+\beta' S^{-1}\beta\sigma_*^2}\right)
\end{aligned}$$

which allows us to get a simple expression for $\hat{\gamma}$ that turns out not to depend on the term $\beta\beta'\sigma_*^2$.

$$\begin{aligned}
\hat{\gamma} &= \left(\beta'\Sigma^{-1}\beta\right)^{-1}\Sigma^{-1}\beta \\
&= \left(\beta' S^{-1}\beta\right)^{-1} S^{-1}\beta \\
&= \beta_1\left(\hat{\beta}\left(\sigma_*^2\right)' S^{-1}\hat{\beta}\left(\sigma_*^2\right)\right)^{-1} S^{-1}\hat{\beta}\left(\sigma_*^2\right)
\end{aligned}$$

and

$$\tilde{\gamma} = \left(\frac{\sigma_*^2}{\beta_1}\right) S^{-1}\hat{\beta}\left(\sigma_*^2\right)$$

Therefore, $\hat{\gamma}$ and $\tilde{\gamma}$ are proportional; the unbiasedness constraint just affects the scale of each vector. Moreover, since $\hat{\beta}\left(\sigma_*^2\right)$ depends on σ_*^2 only through its first argument and S has zeros in its off-diagonal elements on its first row and column, it is clear that $S^{-1}\hat{\beta}\left(\sigma_*^2\right)$ depends on σ_*^2 only through its first argument

$$S^{-1}\hat{\beta}\left(\sigma_*^2\right) = \left(\begin{array}{cc} S_{00}^{-1} & 0 \\ 0 & \hat{S}^{-1} \end{array}\right)\left[\begin{array}{c} \frac{S_{1,0}}{\sigma_*^2} \\ \frac{\hat{C}}{S_{1,0}} \end{array}\right] = \left[\begin{array}{c} S_{00}^{-1}\frac{S_{1,0}}{\beta_1^2\sigma_*^2} \\ \frac{\hat{S}^{-1}\hat{C}}{S_{1,0}} \end{array}\right]$$

This first argument corresponds to the weight on lights in the optimal proxy, and is the only entry of the weight vector that depends on the unknown parameter σ_*^2. Hence, the ratios between any two entries of $S^{-1}\hat{\beta}\left(\sigma_*^2\right)$, and hence of $\hat{\gamma}$ and $\tilde{\gamma}$ that do not correspond to the entry for lights, is pinned down by the data and assumptions A1-A4.

For the analysis in this paper, we will not include the nighttime lights variable as a component of our optimal proxy for true income. We do this because its weight depends on the product $\beta_1^2\sigma_*^2$, which may vary in a range that permits the relative weight on nighttime lights to be zero, or to be infinite.[3] Therefore, without assumptions on β_1 and σ_*^2, we cannot compute this weight, and adding lights does not benefit us in the construction of the proxy.

Finally, the parameters $\alpha_i\left(X\right)$ can be calculated using the system of equations

$$E\left(y_i^n|x_i\right) = \alpha_n\left(x_i\right) + \beta_n E\left(y_i^*|x_i\right)$$

up to the value $E\left(y_i^*|x_i\right)$.

4.2 Calculation of Optimal Forecasts

To calculate absolute magnitudes of $\hat{\gamma}$ (the unbiased estimation weights) and the optimal proxies z_i we need additional assumptions on β_1 and $E\left(y_i^*|X_i\right)$ in order to estimate the $\alpha_i\left(X\right)$'s. Intuitively, the vector $\hat{\gamma}$ incorporates information about cross-country income differences and growth rates, but we need to make assumptions about the average level of our income proxy series. These assumptions are essentially arbitrary but can matter substantially for the results.

We make the following assumption for our baseline analyses:

$$\beta_1 = \left(\hat{\beta}\left(\sigma_*^2\right)' S^{-1}\hat{\beta}\left(\sigma_*^2\right)\right)\left(\sum_{i=1}^{n}\left[S^{-1}\hat{\beta}\left(\sigma_*^2\right)\right]_i\right)^{-1} \text{ and } E\left(y_i^*|X_i\right) = \sum_{n=1}^{N}\gamma_n E\left(y_i^n|X_i\right) \qquad \text{(A5a)}$$

(where $\left(\left[S^{-1}\hat{\beta}\left(\sigma_*^2\right)\right]_i\right)$ is the ith component of $S^{-1}\hat{\beta}\left(\sigma_*^2\right)$). This assumption implies that for the unbiased estimator (with $\hat{\gamma}'\beta = 1$),

$$\eta\left(X_i\right) = 0 \text{ and } \sum_{n=1}^{N}\hat{\gamma}_i = 1$$

so the estimated weights sum to unity and the intercept function of our proxy z_i can be set to zero. HSW

[3] Specifically, if $\phi = \frac{\beta_1^2\sigma_*^2}{S_{11}}$, then $\phi \in \left(\frac{S_{1,0}^2}{S_{00}S_{11}}, \bar{\phi}\right)$, where $\bar{\phi}$ is typically very close to unity. The lower bound sets $\sigma_0^2 = 0$ and assigns infinite relative weight to the lights measure, whereas the upper bound makes the matrix $\hat{S}$ be singular, assigning zero relative weight to the lights measure.

(2012) also consider weights that sum to unity (and that, in fact are also nonzero). The second part of Assumption A5a is motivated by noting that if y_i^n is indexed by time (so $i = (j, t)$ where j indexes countries and t indexes years) and we have

$$\lim_{t\to\infty} y_{j,t}^n = \infty$$

(which is reasonable given that income tends to grow at an exponential rate) then the proxy $z_{j,t}$ will satisfy

$$\lim_{t\to\infty} \frac{z_{j,t}}{\gamma' y_{j,t}} = 1$$

Hence, setting $\eta(X_i) = 0$ will be a good approximation to the value of the optimal proxy for y_i^* in the long run. Since we have no reason to believe that the system governing the relative errors of the national accounts and survey means data is not in a long-run steady state, to which it will eventually tend, we take this normalization as a baseline assumption for computing the optimal proxies.

Another justification for the assumptions on $E(y_i^*|X_i)$ in Assumption A5a is that they yield very similar results to scaling the optimal proxy to national accounts consumption. Bhalla (2002) scales the means of country income distributions to national accounts consumption, arguing that national accounts consumption is an accurate proxy for the fraction of national accounts GDP that is reasonably shared with the poor, and Deaton (2005) also suggests that national accounts consumption may be a reasonable proxy for household disposable income. Most interestingly, the harmonized household disposable income estimates of the Luxembourg Income Study (LIS 2013) seem to confirm this view, coming very close, or much closer than do the surveys used by Chen and Ravallion (2010), to matching national accounts consumption. In fact, LIS household disposable income estimates for OECD countries are virtually identical to World Bank national accounts estimates of consumption per capita in these countries (the average of LIS household disposable income in a dataset of 34 country-years in the OECD on the LIS website is \$24,550, and the same average of their World Bank-recorded consumption is \$24,549). The LIS has very little data on developing countries, but for the 12 country-years in developing countries that they have data for, their estimates of household disposable income are much higher than PovcalNet estimates of mean income or consumption, and for one of these country-years (Guatemala 2006), the LIS estimate even exceeds the World Bank national accounts consumption estimate. A table of these 12 country-years with estimates of mean income or consumption for PovcalNet household surveys, LIS surveys and NA consumption is given as Appendix Table AIII.

We also consider alternative normalizations in which we assume that either national accounts or survey means have a unit relationship with true income (again based on HSW (2012) and CN (2010)) and that the scale of true income matches that of the national accounts or of the survey means:

$$\beta_1 = 1 \text{ and } E\left(y_i^*|X_i\right) = E\left(y_i^1|X_i\right) \text{ (NA)} \tag{A5b}$$

$$\beta_1 = S_{2,0}/S_{1,0} \text{ and } E\left(y_i^*|X_i\right) = E\left(y_i^2|X_i\right) \text{ (Surveys)} \tag{A5c}$$

HSW (2012) also need to assume that

$$\phi = \frac{\beta_1^2\sigma_*^2}{\beta_1^2\sigma_*^2 + \sigma_1^2} \text{ is known.} \tag{A6}$$

in order to compute their estimates. However, since when light is excluded, the unbiased-forecast weights $\hat{\gamma}$ do not depend on σ_*^2 we do not need to make this assumption.

5 Results for Optimal Weights

5.1 Data

We use national accounts data from the World Bank (GDP per capita, PPP, constant 2005 international dollars). The overwhelming majority of countries do not have missing data for this element. National accounts data (from the World Bank or from the Penn World Tables) is overwhelmingly used in cross-country studies of determinants of growth [Barro (1991), Barro and Sala-i-Martin (1992a and b), Mankiw, Romer and Weil (1992), Barro (1999), Sala-i-Martin (1996), Sala-i-Martin, Mulligan and Gil (2002), Sala-i-Martin, Doppelhoffer and Miller (2005), La Porta et al. (1999), Acemoglu et al. (2001, 2002, 2003, 2008), Spolaore and Wacziarg (2005), Ashraf and Galor (2013) among others]

We use a dataset on mean survey consumption from household surveys collected by the World Bank (Povcalnet, http://iresearch.worldbank.org/PovcalNet/index.htm) and used by Chen and Ravallion (2001, 2004, 2010). This dataset mainly consists of surveys after 1990, although there are a few surveys present in the 1980s as well. Many of the survey parameters are heterogeneous (for instance, some surveys are income surveys and others are consumption surveys) but it appears that the heterogeneity is decreasing over time and is not particularly important for our results (allowing indicators for survey income concept does not affect our conclusions). On average, there are about 30-40 surveys each year since 1992, and there are 123 countries surveyed. Survey availability is the primary constraint for our baseline sample from which to estimate the relative optimal weights of national accounts and survey means in the optimal proxy. Overall, we have 701 surveys in this sample, all of which match to national accounts and the lights data for the period 1992-2010.

We also use the lights proxy detailed in Section 3, which we construct by summing up radiance for each country pixel by pixel, dividing by country population and taking the logarithm. This measure is available for all countries after 1992, which precludes using the (relatively few) surveys taken before that date for the analysis. In principle, since we do not use the nighttime lights measure in our index, we could use the surveys from before 1992 for our poverty estimation, but we choose to concentrate on the period 1992-2010 to avoid assumptions about the out-of-sample validity of the optimal survey weights.

Our sample contains observations from the developing world only: there are no World Bank surveys for OECD countries because OECD countries have virtually no population below the $1/day poverty line. Since this paper focuses on poverty, including the OECD countries should not change our analysis. Moreover, lights are a worse measure of output (in particular, growth rates) in OECD countries than in developing countries because the lights measure tends to be topcoded at a light intensity corresponding to the luminosity of a typical developed world city (Doll 2008). Appendix Table AII presents a list of all countries in the base sample, the number and date range of their surveys, and their income as measured by GDP, surveys and lights in the first and last year of their membership in the sample.

For some specifications we also include controls in the function $\alpha(X)$. We have data on log population and the log fractions rural and urban for each country in each year. For most countries and years, we also have data on log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP. All these variables are from the World Bank. Appendix Table AI presents summary statistics for all these variables for the whole world between 1992 and 2010, as well as for the base sample.

It is important to verify explicitly that there indeed exist relationships between nighttime lights, national accounts GDP and true income. To do so, in Table I we provide estimates of the quantity

$$\frac{S_{1,0}}{S_{1,1}S_{0,0}} = \frac{\beta_1\beta_0\sigma_*^2}{S_{1,1}S_{0,0}}$$

which is positive if and only if β_0 and β_1 are both distinct from zero and have the same sign. If β_0 were equal to zero, lights would be a useless indicator of true income and our approach would fail to estimate the structural parameters of our model (we would not be able to identify β_n/β_1 because all covariances would be equal to zero and their quotient would be meaningless). If β_1 were zero or of different sign from β_0, then national accounts would be a useless (or misleading) indicator of true income if lights were considered an indicator of true income and vice versa, making the rationale for using lights to proxy for income questionable. We provide these estimates for a variety of specifications for $\alpha(X_i)$ in rows 1-5: constant $\alpha(X_i)$ (corresponding to looking at the power of lights to predict true income over all across-country-year variation), year fixed effects (looking at how well lights predict cross-country variations in income but not world growth rates), country fixed effects (looking at how well lights predict within-country growth rates but not the cross-country income distribution), country and year fixed effects (looking at cross-country deviations in growth rates), and country and year fixed effects with country time trends (looking at country business cycles). Fortunately, regardless of the type of variation we consider, as Table I shows, lights have predictive power over true income. We also provide estimates of the correlation

$$\frac{S_{2,0}}{S_{2,2}S_{0,0}} = \frac{\beta_2\beta_0\sigma_*^2}{S_{2,2}S_{0,0}}$$

for comparison, which corresponds to the correlation between nighttime lights and survey means. Survey means are also strongly correlated with nighttime lights, though not as strongly as are national account estimates, and cross-country deviations in growth rates of survey means are not correlated with such deviations of nighttime lights (column 4).

5.2 Estimates of Relative Weights

Table II presents estimates of the ratio of the weight of the log survey mean to the weight of the log national accounts mean in the optimal linear proxy for log true income. This ratio corresponds to

$$\hat{\omega} = \hat{\gamma}^2/\hat{\gamma}^1 = \hat{\gamma}_{surveys}/\hat{\gamma}_{NA}$$

or

$$\tilde{\omega} = \tilde{\gamma}^2/\tilde{\gamma}^1 = \tilde{\gamma}_{surveys}/\tilde{\gamma}_{NA}$$

in the notation of Section 4. Recall that this ratio is estimable under Assumptions A1-A4 without any need to assume anything about the magnitude of β_1.

The interesting hypothesis to test on the relative weights that we obtain are not only whether these weights are equal to zero but also how they compare to weights implicitly used in the literature. Research using exclusively national accounts implicitly assumes that $\hat{\gamma}_{surveys} = 0$, and hence that $\hat{\omega} = 0$. Research that exclusively uses survey means implicitly assumes that $\hat{\gamma}_{NA} = 0$ and hence that $\hat{\omega} = +\infty$. Chen and Ravallion (2010) consider a mixed method in which they measure income per capita by the geometric mean of the survey mean consumption and the fitted value of survey mean consumption from a regression of log consumption on a constant and on log consumption in the national accounts. Chen and Ravallion (2010) report that the coefficient on log consumption from the national accounts in such a regression tends to be between 0.6 and 0.85, so we can consider

$$z_i^{CR} = \alpha + \frac{1}{2}y_i^2 + \frac{1}{2}\rho y_i^1$$

where $\rho \in (0.6, 0.85)$. Hence, the Chen-Ravallion (2010) approach assumes that $\hat{\gamma}_{surveys} > \hat{\gamma}_{NA}$, and hence that $\hat{\omega} > 1$.

Each column of Table II presents estimates of optimal relative weights from a different specification of our model. We consider estimates with and without control variables, and we also consider treating each survey as an observation (so countries with more surveys get more weight) or treating each country as an observation (thus countries are equally weighted and surveys in countries with many surveys are underweighted relative to ones in countries with few surveys). Both approaches can be rationalized: surveys within each country differ in methodology, but surveys within a country tend to be more similar than surveys across countries (at least because they are processed by the same statistical offices). In lieu of standard errors we present upper and lower 95% confidence interval bounds for each weight ratio obtained by the bootstrap, which are more conservative than the asymptotic approximation. We also present (as $P(|\hat{\omega}| > 1)$) the fraction of bootstrap iterations in which the weight ratio $\hat{\omega}$ is estimated to be greater than unity in absolute value, which is evidence towards the null hypotheses $\hat{\omega} > 1$ and $\hat{\omega} = +\infty$. We present this statistic because the distribution of $\hat{\omega}$ is nonstandard, and under the null hypothesis $\hat{\omega} = +\infty$ would be bimodal: it would contain no mass in the interval $|\hat{\omega}| > 1$ but a lot of mass outside that interval. The raw confidence interval would then be a misleading indicator of the domain of $\hat{\omega}$ because this domain would no longer be an interval but comprise two disjoint intervals. Hence, the statistic $P(|\hat{\omega}| > 1)$ provides useful information for the few specifications we have with wide confidence intervals for $\hat{\omega}$ by indicating where the mass of the distribution of $\hat{\omega}$ is located.

Each row of Table II presents estimates of this optimal ratio for forecasting different types of variation in true income by including richer and richer fixed effects to account for the remaining variation. Row 1 includes no fixed effects, and thus seeks to find an optimal proxy to capture differences in income across and within countries. Row 2 includes year fixed effects; thus allowing world growth to differ flexibly between national accounts and surveys (this specification also allows differences between satellites in different years). Row 3 includes country fixed effects, thus considering only the relative quality of national accounts and survey means in forecasting country growth rates. Row 4 includes both country and year fixed effects, which focuses on country growth deviations from a world trend. Finally, row 5 includes country and year fixed effects as well as country trends, which looks at the relative quality of national accounts and survey means in forecasting each country's fluctuations around its growth trend, and comparing these fluctuations across countries. Different types of analyses may wish to use relative weights from different columns: for instance, studies of business cycles in developing countries would be most interested in within-country variation around a trend, which is presented in row 5. Since we will be interested in estimating levels of mean income and since we do not wish to impose priors on average levels of income for each country over time, we will be using the estimates from row 1 in our poverty analysis in Section 6.

Our baseline estimate (Row 1 and Column 1 of Table II) suggests that the relative weight of surveys in an optimal proxy, $\hat{\omega}$, is 0.182, and that with 95% confidence, it is between -0.072 and 0.541. Note that we easily reject the null hypothesis that $\hat{\omega} = 1$, or surveys get the same weight as national accounts (Chen and Ravallion 2010), and a fortiriori, $\hat{\omega} = +\infty$. We see that $P(|\hat{\omega}| > 1) < 0.01$ (which is intuitive based on the narrow confidence interval), so virtually all of the distribution of $\hat{\omega}$ is outside the region it would be predicted to be in if surveys had the same weight as national accounts or greater. We also fail to reject the null hypothesis that $\hat{\omega} = 0$, or surveys get zero weight in the optimal proxy [Sala-i-Martin (2002, 2004, 2006), Pinkovskiy and Sala-i-Martin 2009, Sala-i-Martin and Pinkovskiy 2010]. In fact, for any specification in rows 1-4 (using cross-country variation in levels or growth rates) and in any column (with or without control variables for potential common sources of measurement error in nighttime lights, national accounts and survey means) we reject the hypothesis $\hat{\omega} = 1$, we fail to reject the hypothesis $\hat{\omega} = 0$, and we find that $P(|\hat{\omega}| > 1)$ is very low. In row 5, we reject the null hypothesis that $\hat{\omega} = 0$ for two specifications, and we fail to reject the null hypothesis that $\hat{\omega} = 1$ for two other specifications. However, the magnitudes of our estimated relative weights in row 5 are similar to those in rows 1-4 (they are all less than 0.42 in absolute value) and our changes in inference come from differences in the standard errors. In particular, for the specifications with the full set of controls in row 5, the standard errors explode and the estimator distribution has outliers, which prevents us from rejecting the null that $\hat{\omega} = 1$. This is not particularly surprising because we almost succeed in explaining our GDP proxies with the covariate controls as well as with the very rich fixed effects and time trends, leaving little variation to be explained by business cycle variation in true income. It is important to verify that the wide confidence interval comes from large standard errors and not from $\gamma_{NA}/\gamma_{surveys}$ being close to zero; if that were the case, we would see a lot of bootstrap trials with $|\hat{\omega}| > 1$. However, even for these specifications, $P(|\hat{\omega}| > 1)$ is low (it is about 0.15 and 0.06 respectively), suggesting that most of the distribution of $\hat{\omega}$ is still consistent with the hypothesis that the weight on survey means should be lower in absolute value than the weight on the national accounts.

Table III presents estimates of the optimal relative weight of log survey means for each of four large subregions of the developing world (Africa, Latin America, Asia and the post-Communist countries of

Europe and the former USSR) as well as for three time subperiods of the sample (1992-1997, 1998-2003 and 2004-2010). for the same specifications as in Table II. Each cell also presents the estimates of the implied weights of national accounts and of survey means under the additional assumption A5a. We see that for all the subregions and the subperiods, the point estimates of the relative weight of log survey means are less than unity (in fact, less than 0.5), and frequently negative. We also see that the weights on log national accounts GDP are large and close to unity and the weights on log survey means are relatively small, except for certain years. However, we see that there is greater heterogeneity in $\hat{\omega}$ than in Table II and that we can no longer reject most hypotheses of interest based on the bootstrap confidence intervals for our estimates for America, Asia, and the period 2004-2010. This pattern may be rationalized by noting that Latin American countries conduct surveys more frequently than other countries do (and hence, they have had more opportunities to optimize their survey design), that Latin American countries survey income, rather than consumption, and that surveys conducted in the later period 2004-2010 may have been of better quality than preceding ones.

6 Estimates of Global Poverty and True Income Per Capita

6.1 Additional Assumptions on Data for Poverty Estimation

Under assumptions A1-A4 and any one of assumptions A5a-A5c we can calculate the optimal proxies for log true income z_i for each country and year and compute the implied estimates of world poverty. Owing to the paucity of surveys, the literature interpolates or extrapolates survey mean consumption to avoid having poverty estimates depend drastically on whether or not countries with many poor people happen to have a survey in a given year. We perform this imputation by 1) linearly interpolating and extrapolating log survey means for countries with at least two surveys in the Chen-Ravallion database, 2) using the growth rates of national accounts GDP for countries with only one survey in the database, and 3) dropping countries with no surveys in the Chen-Ravallion database.[4] We drop 33 countries this way, of which the largest are South Korea, Afghanistan, Saudi Arabia, Zimbabwe, Cuba, Somalia, the UAE, Libya, Eritrea and Lebanon. Altogether we are left with 123 countries in the developing world, which cover 5.66 billion people in 2010, or about 96.7% of the developing world population.

Having interpolated and extrapolated survey mean consumption, we can easily compute the optimal proxies z_i for the log means of the country income distributions using this interpolated log survey mean series, the log World Bank GDP series, and a set of weights and intercept terms $\eta(X_i)$ from the first row and

[4] Chen and Ravallion (2010) perform a very similar procedure, using national accounts growth rates to interpolate and extrapolate survey means.

column of Table II. We can then use these estimates of the income distribution means to recover poverty by assuming the income distribution is lognormal, recovering the distribution / shape parameters from the Gini coefficients reported with the surveys (which we also interpolate and extrapolate as we do the survey mean consumption for countries with two or more surveys and leave constant for countries with one survey), and integrating up to the poverty line.[5] We follow the World Bank and the United Nations Development Programme and use a poverty line of $1.25 a day in 2005 PPP-adjusted dollars, which is approximately 457 dollars a year. We then bootstrap this procedure for each specification and report the mean, the 5% lower bound and the 95% upper bound of poverty estimates for the years 1992 and 2005 (the first year that lights data are available and the last year of the Chen-Ravallion sample) as well as for each year between 2006 and 2010. The uncertainty in the poverty estimates comes from the fact that the optimal weights and intercept terms used to construct them are estimated with error.[6]

6.2 Estimates of Poverty and True Income per Capita for the Developing World: Baseline Results and Robustness Checks

Table IV presents the poverty rate estimates for the developing world as a whole. Rows 1 and 2 recall the results of the previous literature by presenting poverty estimates under the assumptions that either $\gamma_1 = 0$ and $\gamma_2 = 1$ (designed to replicate the survey mean-based estimates of Chen and Ravallion (2010), hereafter CR (2010)) or, respectively, that $\gamma_1 = 1$ and $\gamma_2 = 0$ (designed to replicate the national account-based estimates of Pinkovskiy and Sala-i-Martin (2009), hereafter PSiM (2009)). Since the interpolation and extrapolation methods are different across papers we cannot replicate the results exactly but we come extremely close. For example, we replicate CR (2010) poverty to be 42% in 1992 and 25.8% in 2005, while in the original paper these numbers are 39.6% in 1993 and 25.2% in 2005 (Row 2). PSiM (2009) estimate poverty to be 8.3% in 1992 and 5.6% in 2005, but these numbers are for the world as a whole rather than for the developing world only, and they also include the countries without surveys. Since it may be safely assumed that no one in rich countries (the OECD) is poor, the population of the OECD is approximately 14% of the world population, and the population of countries without surveys is relatively small, the poverty rates for the developing world implied by PSiM (2009) are 9.5% in 1992 and 6.3% in 2005, while we replicate these rates here to be 9.4% in 1992 and 5% in 2005. The means of the developing world income distributions for these years are reported in Table VI; we see that the national account-based means are more than twice

[5] We use the lognormal distribution as an example, as we have shown in Pinkovskiy and Sala-i-Martin (2010) that neither the interpolation procedures nor the parametric form of the country income distributions matter substantially for estimating the world distribution of income.

[6] Bootstrapping the distribution of our estimator also helps us avoid the problem that we estimate log true income whereas we are interested in estimating true income. We simply use the mean of the distribution of each estimator as our estimate of the desired quantity. In practice, typically, the bias arising from nonlinearity tends to be small, and we would get similar results if we used standard asymptotic analysis.

as large as the survey-based means.

The rest of Table IV presents our new estimates of developing world poverty based on optimally combining national accounts and survey means. Row 3 presents our baseline estimates under the long-run scaling assumption A5a. We see that our poverty estimate for 1992 is 11.8%, which falls to 6.1% in 2005 and 4.5% in 2010. Our estimated poverty rates are very close to the estimates of PSiM (2009), and we can reject with 95% confidence the hypothesis that poverty fell by less than half by 2010. Hence (and not surprisingly given our evidence on relative weights in Table II) optimally combining national accounts and survey means through the use of the nighttime lights data as an independent benchmark to uncover the joint relationship of these measures' errors from true income yields poverty estimates much closer to those deriving from the national accounts than from the survey means. Table VI presents our baseline estimates of developing world true income per capita levels; they are much closer to the national account-based means than to the survey-based means.

Rows 4 and 5 present robustness checks of this result by changing assumption A5a to assumptions A5b and A5c respectively; hence, by normalizing the optimal proxy to the level of the national accounts or to that of the household surveys. We see that normalization makes a difference: the level of poverty that we calculate under assumption A5c (normalizing to surveys) is much higher than the one that we calculate under assumption A5b (normalizing to national accounts). However, even under assumption A5c, which uses very nearly the same scale for income as do CR (2010) and uses the weights only to compute growth rates and cross-sectional differences across countries, we see that poverty is estimated to be a third to a half the size in all years considered than in CR (2010), and that our survey-normalized estimates indicate both lower and faster-falling poverty rates than do the estimates of CR (2010) with 95% confidence. Table VI presents the corresponding estimates of the developing world true income per capita levels; they vary as the poverty rate estimates do. Figures III and IV present the time paths of world poverty rates, the first in levels and the second as a percentage of the 1992 value. We see that even when we use the survey normalization, poverty estimated using the optimal weighting method is much lower and falls faster than poverty estimated using surveys alone.

Rows 6 and 7 again use assumption A5a to scale our optimal proxies, and explore the sensitivity of our results to assumption A1: the homogeneity of the underlying statistical model across countries and years. Row 6 presents estimates for which the relative weights have been re-estimated in each year using a sample of countries and years with surveys in that year only. This check is important because surveys may be improving or deteriorating over time; also, satellites in different years may have different optical properties and record the same lights differently. To avoid sharp changes in poverty estimates when weights change from year to year, we normalize these estimates using a recursive formula. [7] Since in Row 7 we allow the weights to vary cross-sectionally rather than longitudinally, no changes to the scaling assumption are required. We see that the poverty estimates are again quite similar to the baseline, albeit with wider confidence intervals. Rows 8 and 9 add covariates to our baseline specification in order to account for factors that may bias light density away from true income, such as the shares of manufacturing or services in GDP or the fraction of the population that is urban. Once again, our poverty estimates are very similar to the baseline. Table VI presents the corresponding estimates of the developing world true income per capita levels.[8]

Our last two robustness checks attempt to address conceptual problems with combining national accounts and survey means. It is obvious (Deaton 2005, 2010) that national accounts GDP and survey mean consumption measure two different income concepts, and that a part of their divergence is explained by this difference. It may be the case that while true income is growing faster than survey consumption, it is not necessarily reflected by national accounts GDP and is in fact much closer in concept to national accounts consumption. However, since we do not include national accounts consumption as a potential component of our optimal proxy in our baseline specification, we cannot rule this out. Row 10 adds log national accounts consumption (household final consumption expenditure) as another proxy in our optimal construction of true income, hence as y_i^3 in the notation of Section 4. Hence, we now compute our proxy as a weighted average of national accounts GDP, survey means and national accounts consumption and we have three weights. Additionally, instead of using assumption A5a to scale our optimal proxy, we instead scale the proxy to the level of national accounts consumption: we assume that $E\left(y_i^*|X_i\right) = E\left(y_i^3|X_i\right)$, and that $\beta_3 = 1$.[9] Our

[7] The assumption that

$$E\left(y_{i,t}^*|x_{i,t}\right) = \gamma_{1,t} E\left(y_{i,t}^1|x_{i,t}\right) + \gamma_{2,t} E\left(y_{i,t}^2|x_{i,t}\right)$$

is modified to read

$$E\left(y_{i,t}^*|x_{i,t}\right) = \lambda_{1,t} E\left(y_{i,t}^1|x_{i,t}\right) + \lambda_{2,t} E\left(y_{i,t}^2|x_{i,t}\right)$$

where

$$\lambda_{1,t+1} = (1-g)\,\lambda_{1,t} + g\gamma_{1,t}$$

$$\lambda_{2,t+1} = (1-g)\,\lambda_{2,t} + g\gamma_{2,t}$$

$$g = \lambda_{1,t}\left(E\left(y_{i,t+1}^1|x_{i,t+1}\right) - E\left(y_{i,t}^1|x_{i,t}\right)\right) + \lambda_{2,t}\left(E\left(y_{i,t+1}^2|x_{i,t+1}\right) - E\left(y_{i,t}^2|x_{i,t}\right)\right)$$

and the initial values of $\lambda_{1,t}$ and $\lambda_{2,t}$ are set to the baseline (Row 1) values of γ_1 and γ_2

[8] These estimates of the true income per capita levels may seem somewhat puzzling because 1) the year-specific weights estimates are lower than the baseline estimates, although the poverty levels estimated with year-specific weights are also lower than the baseline poverty levels, and 2) the region-specific weights estimates of per capita true income appear to be implausibly large, although the poverty rates are not out of line with our other poverty estimates. However, these apparent puzzles can be easily reconciled by noting that poverty is a nonlinear function of true income and therefore, that the lower tail of the distribution of our estimates for true income matters more than does the uppore tail of that distribution, and potentially, more than does the mean. We see that in fact, the lower 5% percentile of the distribution of the year-specific-weights true income estimates tends to be higher than the lower 5% percentile of the distribution of the baseline true income estimates, and that the lower 5% percentile of the distribution of the region-specific-weights true income estimates is very similar to that of the baseline true income estimates. The large positive outliers for the region-specific-weights true income estimates matter little for poverty estimation because they replace the top percentiles of the baseline estimates distribution, which also generate very low poverty rates (since poverty rates are bounded below by zero).

[9] Using Assumption A5a yields poverty estimates that are virtually identical to the baseline estimates.

poverty estimates hardly change.[10] This robustness check is also notable because it shows that assumption A5a is very similar to just scaling the optimal proxy to consumption, which appears to be consistent with survey evidence from the Luxembourg Income Study (LIS 2013).

A second problem is that if household surveys systematically mismeasure the mean of the income distribution, they may also systematically mismeasure its dispersion. There is no reason, either based on theory or on data, to believe that the household surveys understate income inequality, and in our case there are good reasons to believe that they actually overstate it. Deaton (2005) presents a parametric example in which nonresponse leads surveys to underestimate the mean but not inequality. Survey estimates of disposable income from the Luxembourg Income Study (LIS 2013) find mean incomes to be larger and Gini coefficients to be smaller for the several developing countries and years for which both LIS estimates and survey estimates from Chen and Ravallion (2010) are available.[11]. Finally, as discussed in Section 2, it is very unlikely that any supplementary income indicated by the nighttime lights data (arising from proper valuation of housing and public goods) is particularly unequally distributed as this income is embodied in capital-intensive and bulky goods that are unlikely to be very closely held. However, in a final robusteness check we will explicitly consider the possibility that the gains from global growth have been increasingly distributed to rich people in developing countries who do not cooperate with household surveys. Such a process would cause us to attribute to the poor the income that went to the rich and to overestimate poverty reduction. To assess the sensitivity of our results to mismeasurement of inequality, we consider the situation in which the Gini coefficient for each survey is downward biased by a certain number of Gini points. We select this number by computing the median within-country standard deviation of the Gini coefficient under a weighted scheme in which each country is weighted by the standard deviation of the years in which surveys for that country are available (hence, countries with surveys over a longer period of time get more weight), and by taking 1.96 times this number, which amounts to 5.37 Gini points. (This mismeasurement is somewhat below the median within-country range of the Gini coefficient (computed in the same manner), which is 7.6 Gini points.) We compute poverty rates using the new inequality series in which 5.37 is added to each country's Gini coefficient. We also compute an additional conservative estimate of the ratio of poverty in 2010 to poverty in 1992 by using the poverty based on the new inequality data for 2010 and poverty based on the actual CR (2010) inequality data for 1992. Such a computation would account for the possibility that surveys become progressively worse at measuring inequality, perhaps because much growth accrues to rich

[10] This is not surprising because the weight on the surveys remains small (the ratio to the national accounts GDP weight is 0.179). The ratio of the weight on national accounts consumption to national accounts GDP is smaller (0.113) but imprecisely estimated. Hence, the weights are about 77% national accounts GDP, 14% surveys, and 9% national accounts consumption.

[11] For example, the LIS survey for Brazil finds that mean disposable income is $6000 and the Gini is 48; the Brazilian survey cited by Chen and Ravallion (2010) finds that mean income is $3900 and the Gini is 56. Comparisons for Colombia, Estonia, Guatemala, Hungary, Mexico, Peru and Poland are similar.

people who do not participate in surveys, and this procedure would make it more difficult to find poverty declines. Row 11 presents the resulting estimates; under the new inequality series, poverty is higher and falls more slowly, but is much closer to the baseline than to the survey-based poverty series, let alone the CR (2010) estimates. We also find that even under our conservative procedure for computing poverty ratios, by 2010, poverty has fallen to less than 57% of its 1992 level. Table VI presents the corresponding estimates of the developing world true income per capita levels.

The conclusion that we draw from Tables IV and Table VI is that the developing world has grown by much more, and poverty has fallen by much more than indicated by the household surveys alone. For all specifications, even the ones in which the overall scale of our true income measure is set to be the same as that of the household surveys, poverty in 2010 (and in all other years) is estimated to be statistically significantly lower with our optimal weighting method than by using survey means alone. The difference is also practically large: our largest estimate for poverty in 2010 is 12.1%, as compared with 20.5% using only survey means. For all specifications except the conservative ratio computation with measurement error in inequality, we find that poverty declined by a larger percentage from 1992 to 2010 than we would find based on evidence from household surveys, and that this decline happened off of a lower poverty baseline. We find that the most important factors affecting our estimates are our choice of scaling of the true income measure, and to a lesser extent, our assumptions about mismeasurement of inequality by the surveys. It is intuitive that these two factors should be the most important, as, in principle, assuming very low levels of true income or assuming that all the income growth since 1992 went to the nonpoor would be enough to remove any poverty decline whatsoever. However, these assumptions either imply that the data we have are systematically untrustworthy, or imply that we are away from the long-run steady state of the process governing the evolution of national accounts and survey means.

6.3 Regional Results

Tables V and VII present poverty and true income per capita estimates for various regions of the developing world. Each row reports a different specification, which are the same specifications as in Tables IV and VI. We report only a few poverty numbers for each region in order to present a compact picture, and we only present the 2010 / 1992 poverty ratio upper confidence bound as a tool for inference. In Table VII we just report true income levels for 1992, 2005 and 2010 without confidence intervals. We see much the same pattern as for the world as a whole, with East and South Asia experiencing more rapid poverty reduction and Sub-Saharan Africa experiencing less rapid poverty reduction. Interestingly, for all specifications except the one using the survey mean normalization (row 5) and the conservative ratio using the increased inequality series (last cell of row 11), Sub-Saharan Africa reduces poverty by more than 30% between 1992 and 2010, which is statistically significantly different from the 20% reduction one obtains by just using survey means (row 2 of the table). Since both of these robustness checks are somewhat extreme (given that national accounts get a much higher weight than surveys do, it is not particularly plausible that the scale of the optimal true mean proxy should be so far away from its long-run value as are the survey means; it is neither likely that African surveys drastically undermeasured inequality in 2010 relative to 1992), this suggests that Africa is doing better than is suggested by the evidence in the household surveys.[12]

7 Conclusion

A large number of papers have attempted to estimate poverty rates around the globe. All of them use survey data to determine the dispersion of income across citizens around a given mean to construct the distribution of income of each country and then they estimate the poverty rates as the integral of that distribution to the left of a given poverty line. Different papers use different types of surveys, different methods to parameterize each country's distribution of income, different ways to interpolate and extrapolate with missing observations, different data sources and different estimates of the mean of each country distribution of income or consumption. Our reading of the literature is that the final estimates of the global poverty rate do not depend crucially on the exact parametric specifications chosen by the researchers nor do they depend on the way they interpolate or extrapolate the missing data (Pinkovskiy and Sala-i-Martin 2009, Dhongde and Minoiu 2010). The determining methodological choice is the choice of anchor of the income distribution. In this sense, there are two groups of papers. There are those that anchor the distribution of income to the national accounts' GDP per capita [Bhalla (2002), Sala-i-Martin (2002, 2004, 2006), Pinkovskiy and Sala-i-Martin (2009) and Sala-i-Martin and Pinkovskiy (2010)]. And then there are those that anchor the distribution to the survey means [Chen and Ravallion (2001, 2004, 2010), Milanovic (2005)]. The choice of the mean of the distribution matters empirically because it turns out that, for many developing countries, the survey means not only are much smaller than the national accounts' GDP per capita, but they also

[12] Sala-i-Martin and Pinkovskiy (2010) use exclusively national accounts to conclude that Africa is on track to achieve the Millennium Development Goal of halving poverty relative to the 1990 level by 2015. Our poverty ratios compare poverty in 2010 and 1992 only and therefore cannot be used to answer this question; if we forecast poverty to 2015 and compute a 2015/1992 poverty ratio we would get that Africa reduces poverty by 2015 to 55% or less of its 1992 level for all specifications except for the two specifications mentioned in this paragraph. Our estimates of the 2010/1992 Africa poverty ratio using our baseline weights are higher than using national accounts alone (row 2) because the baseline estimates place some positive weight on survey means, and we know that the growth rate of surveys is smaller than the growth rate of GDP. However, it is likely that this weight on the survey means is too large in the context of Africa. From Table III we see that the weight on surveys for the African subsample alone is negative, in contrast with the small positive weight on surveys for the whole world sample. In row 7 of Table V we estimate the 2010/1992 African poverty ratio using the weights estimated off of the African subsample only, and we see that this ratio is actually lower than the ratio we obtain using national accounts alone.

grow much more slowly. Obviously, if one anchors the distribution to a smaller number, one obtains a much larger poverty rate. And if the anchor grows at a smaller speed, the poverty rate will decline much more slowly. Hence, the studies that use the estimated average income of the survey as the mean of each country's distribution tend to find much larger poverty rates than the studies that use per capita GDP. And they also tend to estimate that these poverty rates fall much more slowly.

Nobody really knows why the survey means are different from the estimates of GDP per capita (Deaton 2005). The rationale for using the national account's GDP per capita is that the distribution of income should be consistent with all the macroeconomic studies used to evaluate the performance of countries. When economists say that China grew at $x\%$ per year during an entire decade, what they mean is that its GDP per capita (not its survey means) grew at $x\%$ per year. And when they put the growth rate for China in a cross country comparison analysis, they use the growth rate of GDP per capita. And any measure of the distribution of income should be consistent with the most widely used measure of income: GDP. If the survey means are smaller than GDP per capita, it must be due to some kind of misreporting on the part of the surveyed. Economists using GDP as the anchor implicitly assume that the missing income occurs proportionally across the entire income distribution.

Researchers that like to use the survey mean, on the other hand, argue that it is possible that much of the income missing from the surveys goes to the nonpoor (Chen and Ravallion 2010). Hence, even though GDP is a good measure of overall income, when it comes to estimating poverty the survey means are much closer to the mean of the "distribution of the poor". Since nobody knows for sure the source of the discrepancy between GDP per capita and the survey means, we cannot be sure whose estimates of poverty rates are more accurate.

We believe that this paper provides an avenue to solve the problem. We use a third, independently collected data on economic activity to test whether GDP per capita or survey means are a better estimate of true income. The data we use is satellite-recorded luminosity at night as measured by the DMSP-OLS satellites of the National Oceanic and Atmospheric Administration (NOAA).

In general, a positive correlation between measured income (national accounts or survey means) and nighttime lights could be due to two factors: that they are both correlated with true income, or that their measurement errors are strongly correlated with each other. However, the latter possibility is implausible because the generating process of nighttime lights data is to a very large degree independent of the generating process either of national accounts or of survey means. For example, measured income is collected by statisticians interacting with survey respondents, while nighttime lights are recorded impersonally by satellites. Statistical teams use different procedures in different countries, while lights are recorded homogeneously across national borders. Both national accounts and survey means may suffer from nonrandom nonresponse and misreporting, whereas nighttime lights do not require compliance or truthfulness of the surveyed population to record whatever lights exist. Moreover, nighttime lights may vary because of climatic conditions such as auroral activity, cloudiness and humidity, or because of cultural attitudes towards lighting, which presumably do not affect measurement errors in national accounts or survey means. Therefore, the strength of the correlation between nighttime lights and measured income is directly related to the strength of the correlation between the given income measurement and the true income it is trying to measure. We can use the ratios of correlations between nighttime lights and different income measurements to assess the relative strengths of the correlations between these income measurements and unobserved true income.

Using data on nighttime lights we test whether national accounts or survey means better reflect variation in true income across countries and over time. We find that national accounts do a better job. We also use the luminosity data to create a new proxy for true income as a log linear weighted average of the national accounts and the survey means. We find that the weight that we wish to place on survey means is 18% of the weight that we wish to place on national accounts GDP.

Finally, we use the new optimal measure of true income to calculate the evolution of poverty at the worldwide level as well as at the regional level. Not surprisingly, our estimates of poverty rates are between those of the literature that uses GDP and the literature that uses survey means. Given that our optimal measure gives a small weight to survey means, our optimal estimates of poverty rates tend to be closer to those reported in the research that uses GDP as the anchor. An objection could be that surveys not only mismeasure the mean of the distribution of income, but also inequality, and that it is therefore incorrect to combine survey-based inequality measures with income distribution means that are constructed on the basis of national accounts. However, we show that poverty declines more rapidly if measured using our optimally constructed means than if measured using survey means alone even if we allow for very large survey errors in inequality measurement.

And this is the main conclusion of this paper: poverty rates have been falling much faster than predicted by the literature that measures poverty solely using survey means.

References

[1] **Acemoglu, Daron, Simon Johnson, and James A. Robinson**. 2001. "The Colonial Origins of Comparative Development: An Empirical Investigation." *American Economic Review* 91, no. 5: 1369-1401.

[2] **Acemoglu, Daron, Simon Johnson, and James A. Robinson**. 2002. "Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution." *Quarterly Journal of Economics* 117, no. 4: 1231-1294.

[3] **Acemoglu, Daron, Simon Johnson, and James Robinson**. 2003. "Disease and Development in Historical Perspective." *Journal of The European Economic Association* 1, no. 2-3: 397-405.

[4] **Acemoglu, Daron, Simon Johnson, James A. Robinson, and Pierre Yared**. 2008. "Income and Democracy." *American Economic Review* 98, no. 3: 808-842.

[5] **Alesina, Alberto F., Stelios Michalopoulos, and Elias Papaioannou**. 2012. "Ethnic Inequality." NBER WP #18512.

[6] **Alvaredo, Facundo and Leonardo Gasparini**. 2013. "Recent Trends in Inequality and Poverty in Developing Countries." Documiento de Trabajo #151, Universidad de la Plata. Forthcoming *Handbook of Income Distribution*, vol. 2, eds. Atkinson, Anthony B. and Francois Bourguignon.

[7] **Ashraf, Quamrul, and Oded Galor**. 2013. "The 'Out of Africa' Hypothesis, Human Genetic Diversity, and Comparative Economic Development." *American Economic Review* 103, no. 1: 1-46.

[8] **Barro, Robert J.** 1991. "Economic Growth in a Cross Section of Countries." *Quarterly Journal of Economics* 106, no. 2: 407-443.

[9] **Barro, Robert J.** 1999. "Human Capital and Growth in Cross-Country Regressions." *Swedish Economic Policy Review* 6, no. 2: 237-277.

[10] **Barro, Robert J., and Xavier Sala-i-Martin**. 1992a. "Convergence." *Journal of Political Economy* 100, no. 2: 223-251.

[11] **Barro, Robert J., and Xavier Sala-i-Martin**. 1992b. "Public Finance in Models of Economic Growth." *Review of Economic Studies* 59, no. 4: 645-661.

[12] **Bhalla, Surjit S.** 2002. *Imagine there's no country: Poverty, inequality, and growth in the era of globalization*. Washington, D.C.:Penguin.

[13] **Bourguignon, Francois, and Christian Morrisson**. 2002. "Inequality among World Citizens: 1820-1992." *American Economic Review* 92, no. 4: 727-744.

[14] **Chen, Shaohua, and Martin Ravallion**. 2001. "How Did the World's Poorest Fare in the 1990s?." *Review of Income and Wealth* 47, no. 3: 283-300.

[15] **Chen, Shaohua, and Martin Ravallion**. 2010. "The Developing World Is Poorer Than We Thought, but No Less Successful in the Fight against Poverty." *Quarterly Journal of Economics* 125, no. 4: 1577-1625.

[16] **Chen, Shaohua, Martin Ravallion and Prem Sangraula**. 2013. *Povcalnet*, http://iresearch.worldbank.org/PovcalNet/index.htm. Accessed October 1, 2013.

[17] **Chen, Shaohua, and Martin Ravallion**. 2004. "How Have the World's Poorest Fared since the Early 1980s?." *World Bank Research Observer* 19, no. 2: 141-169.

[18] **Chen, Xi, and William D. Nordhaus**. 2010. "The Value of Luminosity Data as a Proxy for Economic Statistics." NBER Working Papers: 16317.

[19] **Deaton, Angus**. 2001. "Counting the World's Poor: Problems and Possible Solutions." *World Bank Research Observer* 16, no. 2: 125-147.

[20] **Deaton, Angus**. 2005. "Measuring Poverty in a Growing World (or Measuring Growth in a Poor World)." *Review of Economics and Statistics* 87, no. 1: 1-19.

[21] **Deaton, Angus**. 2010. "Price Indexes, Inequality, and the Measurement of World Poverty." *American Economic Review* 100, no. 1: 5-34.

[22] **Dhongde, Shatakshee and Camelia Minoiu**. 2010. "Global Poverty Estimates: Present and Future." ECINE Working Paper Series, ECINEQ-WP 2010 #281

[23] **Doll, Christopher N.H.** 2008. "CIESIN Thematic Guide to Night-time Light Remote Sensing and its Applications." Manuscript.

[24] **Doll, Christopher N.H.**, **Jan Peter Muller and Jeremy G. Morley**. 2006. "Mapping Regional Economic Activity from Night-Time Light Satellite Imagery." *Ecological Economics* 57: 75-92

[25] **Elvidge, Christopher D. and Kimberly E. Baugh, Eric A. Kihn, Herbert W. Kroehl, Ethan R. Davis**. 1997. "Mapping City Lights With Nighttime Data from the DMSP Operational Linescan System." *Photogrammetric Engineering & Remote Sensing* 63, no. 6: 727-734

[26] **Elvidge, Christopher D., Kimberly E. Baugh, John B. Dietz, Theodore Bland, Paul C. Sutton and Herbert W. Kroehl.** 1999. "Radiance Calibration of DMSP-OLS Low-Light Imaging Data of Human Settlements." *Remote Sensing of Environment*, 68(1):77-88.

[27] **Elvidge, Christopher D., Kimberly E. Baugh, Sharolyn. J. Anderson, Paul. C. Sutton, and Tilottama Ghosh**. 2012. "The Night Light Development Index (NLDI): A Spatially Explicit Measure of Human Development from Satellite Data." *Social Geography* 7, 23–35.

[28] **Ghosh, Tilottama, Rebecca L. Powell, Christopher D. Elvidge, Kimberly E. Baugh, Paul C. Sutton and Sharolyn Anderson**. 2010. "Shedding Light on the Global Distribution of Economic Activity." *The Open Geography Journal* 3, 148-161.

[29] **Henderson, J. Vernon, Adam Storeygard, and David N. Weil**. 2009. "Measuring Economic Growth from Outer Space." NBER Working Paper #15199,

[30] **Henderson, J. Vernon, Adam Storeygard, and David N. Weil**. 2012. "Measuring Economic Growth from Outer Space." *American Economic Review* 102, no. 2: 994-1028.

[31] **Johnson, Simon, William Larson, Chris Papageorgiou, and Arvind Subramanian**. 2013. "Is Newer Better? Penn World Table Revisions and Their Impact on Growth Estimates." *Journal of Monetary Economics* 60, no. 2: 255-274.

[32] **Karshenas, Massoud**, 2003, "Global Poverty: National Accounts Based versus Survey Based Estimates," *Development and Change* 34(4): 683-712.

[33] **Baugh, Kimberly E., Christopher D. Elvidge, Tilottama Ghosh and Daniel Ziskin**. 2009. "Development of a 2009 Stable Lights Product using DMSP-OLS Data." *Proceedings of the 30th Asia-Pacific Advanced Network Meeting.*

[34] **Korinek, Anton, Johan Mistiaen and Martin Ravallion**, 2006. "Survey Nonresponse and the Distribution of Income." *Journal of Economic Inequality*, 4(2): 33–55.

[35] **La Porta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer and Robert Vishny**. 1998. "Law and Finance." *Journal of Political Economy* 106, no. 6: 1113-1155.

[36] **La Porta, Rafael, Florencio Lopez-de-Silanes, and Andrei Shleifer**. 2008. "The Economic Consequences of Legal Origins." *Journal of Economic Literature* 46, no. 2: 285-332.

[37] **Luxembourg Income Study (LIS) Database, Key Figures**, http://www.lisdatacenter.org (multiple countries; [1967-2010]). Luxembourg: LIS.

[38] **Mankiw, N. Gregory, David Romer, and David N. Weil**. 1992. "A Contribution to the Empirics of Economic Growth." *Quarterly Journal of Economics* 107, no. 2: 407-437.

[39] **Michalopoulos, Stelios, and Elias Papaioannou**. 2013. "Pre-colonial Ethnic Institutions and Contemporary African Development." *Econometrica* 81, no. 1: 113-152.

[40] **Michalopoulos, Stelios, and Elias Papaioannou**. 2014. "National Institutions and Sub-national Development in Africa." *Quarterly Journal of Economics*, forthcoming.

[41] **Milanovic, Branko**. *Worlds Apart: Measuring International and Global Inequality*. Princeton and Oxford: Princeton University Press. 2005.

[42] **Mulligan, Casey B., Ricard Gil, and Xavier Sala-i-Martin**. 2004. "Do Democracies Have Different Public Policies than Nondemocracies?." *Journal of Economic Perspectives* 18, no. 1: 51-74.

[43] **Pinkovskiy, Maxim L.** 2013. "Economic Discontinuities at Borders: Estimation Using Satellite Data of Lights at Night." mimeo, MIT.

[44] **Pinkovskiy, Maxim L**. 2013. "World Welfare Is Rising: Estimation Using Nonparametric Bounds on Welfare Measures." *Journal of Public Economics*, v. 97, pp. 176-95

[45] **Pinkovskiy, Maxim L. and Xavier Sala-i-Martin**. 2009. "Parametric Estimations of the World Distribution of Income." NBER Working Paper #15433.

[46] **Sala-i-Martin, Xavier**. 1997. "I Just Ran Two Million Regressions." *American Economic Review* 87, no. 2: 178-183.

[47] **Sala-i-Martin, Xavier**. 2002a. "The Disturbing 'Rise' of Global Income Inequality." NBER Working Paper #8904.

[48] **Sala-i-Martin, Xavier**. 2002b. "The World Distribution of Income (estimated from Individual Country Distributions)." NBER Working Paper #8933.

[49] **Sala-i-Martin, Xavier**. 2006. "The World Distribution of Income: Falling Poverty and . . . Convergence, Period." *Quarterly Journal of Economics* 121, no. 2: 351-397.

[50] **Sala-i-Martin, Xavier, Gernot Doppelhofer, and Ronald I. Miller**. 2004. "Determinants of Long-Term Growth: A Bayesian Averaging of Classical Estimates (BACE) Approach." *American Economic Review* 94, no. 4: 813-835.

[51] **Sala-i-Martin, Xavier and Maxim L. Pinkovskiy** 2010. "African Poverty is Falling...Much Faster Than You Think!" NBER Working Paper #15775.

[52] **Spolaore, Enrico and Romain Wacziarg**. 2009. "The Diffusion of Development," *Quarterly Journal of Economics*, 124 no.2: pp 469-592

[53] **Sutton, Paul C., Christopher D. Elvidge and Tilottama Ghosh**. 2007. "Estimation of Gross Domestic Product at Sub-National Scales using Nighttime Satellite Imagery." *International Journal of Ecological Economics & Statistics* 8:5-21

[54] **World Bank**. 2013. *World Development Indicators*. Washington, DC.

[55] **Young, Alwyn**. (2012). "The African Growth Miracle." *Journal of Political Economy* 120, no. 4, pp. 696-739.

8 Tables

Table I (I)

Correlations Between Measured Income and Nighttime Lights		
	(1)	(2)
Specification	National Accounts	Survey Means
No FE	.862 (.809) (.904)	.799 (.743) (.852)
Year FE	.870 (.817) (.911)	.809 (.753) (.859)
Country FE	.461 (.336) (.583)	.275 (.172) (.383)
Country FE + Year FE	.450 (.309) (.584)	.104 (-.068) (.229)
Ctry FE + Ctry Trend + Year FE	.338 (.144) (.550)	.233 (.112) (.365)

The first column of Table I presents estimates and 95% confidence intervals for the correlation between log nighttime lights per capita and log national accounts GDP per capita, and the second column presents estimates and 95% confidence intervals for the correlation between log nighttime lights per capita and log survey mean income or consumption per capita, as described in Section 5. Correlations estimated over sample of countries and years with available survey and nighttime lights data, with 701 observations and 123 different countries. Data on nighttime lights from the NOAA, data on national accounts GDP from the World Development Indicators, and data on survey means is from Chen and Ravallion (2010). Confidence intervals obtained by boostrapping countries with 120 repetitions. Each row corresponds to partialling out different fixed effects, mentioned in the row headings, from the log income measures (including them in the function $\alpha(X_i)$).

Table II

(II)

Estimates of Relative Weight of Survey Means in Optimal GDP Index						
	(1)	(2)	(3)	(4)	(5)	(6)
Specification	Base Line	Equal. Weight	Urb / Rur Controls	Urb / Rur Eq. Wt.	All Controls	All Ctrls. Eq. Wt.
No FE	**.182**	-.020	.111	-.024	.079	-.097
Confidence Bounds	**(-.072)**	(-.220)	(-.158)	(-.249)	(-.247)	(-.375)
	(.541)	(.246)	(.525)	(.373)	(.486)	(.224)
P-value $\lvert\hat{\omega}\rvert > 1$	**(.008)**	(0)	(.008)	(0)	(.008)	(0)
Year FE	.221	.002	.156	-.009	.130	-.051
Confidence Bounds	(-.048)	(-.231)	(-.122)	(-.247)	(-.155)	(-.322)
	(.612)	(.273)	(.523)	(.379)	(.600)	(.255)
P-value $\lvert\hat{\omega}\rvert > 1$	(0)	(0)	(.008)	(0)	(.008)	(0)
Country FE	-.052	-.036	-.052	-.002	-.080	-.122
Confidence Bounds	(-.274)	(-.321)	(-.313)	(-.331)	(-.328)	(-.348)
	(.324)	(.355)	(.390)	(.442)	(.249)	(.131)
P-value $\lvert\hat{\omega}\rvert > 1$	(0)	(0)	(0)	(0)	(0)	(0)
Country FE + Year FE	-.036	-.025	-.030	-.005	-.024	-.061
Confidence Bounds	(-.240)	(-.266)	(-.246)	(-.275)	(-.235)	(-.286)
	(.299)	(.240)	(.351)	(.284)	(.340)	(.231)
P-value $\lvert\hat{\omega}\rvert > 1$	(0)	(0)	(0)	(0)	(0)	(0)
Ctry FE + Ctry Trend + Year FE	.191	.168	.169	.162	.413	-.224
Confidence Bounds	(.019)	(.022)	(-.002)	(-.018)	(-3.862)	(-.089)
	(.460)	(.433)	(.509)	(.475)	(3.732)	(2.546)
P-value $\lvert\hat{\omega}\rvert > 1$	(0)	(0)	(0)	(0)	(.15)	(.066)
No. Obs.	701	701	701	701	650	650
No. Clusters	123	123	123	123	117	117

Each column of Table II presents estimates and 95% confidence intervals for $\hat{\omega} = \hat{\gamma}_{surveys}/\hat{\gamma}_{NA}$, the ratio of the weight of log survey means per capita to the weight of log national accounts GDP per capita in the optimal unbiased proxy z_i of the mean of the true income distribution. The different column specifications involve different sample survey weighting schemes, control variables, and inclusion of national accounts consumption in z_i. The baseline specification weighs all surveys equally, does not include covariate controls, and proxies z_i by a combination of national accounts GDP per capita and survey means per capita only. The bolded cell indicates the estimates that will be used in the baseline specifications in all subsequent tables. Data definitions, inference procedures and sample selection are as in Table I. Columns 1, 3, and 5 correspond to weighting all observations equally; columns 2, 4, 6, 8 and 10 correspond to weighting all countries equally. The controls in columns 3 and 4 are: log total population, log percentage rural population, log percentage urban population. The controls in columns 5 and 6 are: log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP. Each row corresponds to partialling out different fixed effects, mentioned at the foot of the table, from the log income measures (including them in the function $\alpha(X_i)$).

Table III (III)

Estimates of Weights on Survey Means and National Accounts								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Specification	Base Line	Equal. Weight	Urb / Rur Controls	Urb / Rur Eq. Wt.	All Controls	All Ctrls. Eq. Wt.	Obs.	Clusters
Baseline	**.182**	-.020	.111	-.024	.079	-.097	701	123
	(-.072)	(-.220)	(-.158)	(-.249)	(-.247)	(-.375)		
	(.541)	(.246)	(.525)	(.373)	(.486)	(.224)		
Survey Weight	**.154**	-.021	.100	-.024	.073	-.107		
WB GDP Weight	**.845**	1.021	.899	1.024	.926	1.107		
Region == Africa	-.210	-.224	-.176	-.175	-.458	-.457	114	41
	(-.560)	(-.526)	(-.578)	(-.536)	(-7.417)	(-2.845)		
	(.219)	(.222)	(.360)	(.519)	(1.395)	(1.915)		
Survey Weight	-.266	-.289	-.214	-.213	-.845	-.842		
WB GDP Weight	1.266	1.289	1.214	1.213	1.845	1.842		
Region == Asia	.077	-.175	-.017	-.307	-.137	-.224	119	29
	(-.823)	(-.829)	(-1.221)	(-1.171)	(-4.186)	(-4.572)		
	(2.654)	(1.006)	(2.178)	(1.540)	(1.418)	(1.791)		
Survey Weight	.071	-.212	-.017	-.444	-.159	-.289		
WB GDP Weight	.928	1.212	1.017	1.444	1.159	1.289		
Region == America	.441	.218	-.529	.394	.160	.082	234	25
	(-.328)	(-.291)	(-.378)	(-.504)	(-.393)	(-.392)		
	(2.678)	(1.332)	(1.901)	(2.356)	(.805)	(.880)		
Survey Weight	.306	.179	-1.124	.283	.138	.076		
WB GDP Weight	.693	.820	2.124	.716	.861	.923		
Region == PostCommunist	.145	-.014	.095	.025	-.082	-.209	234	28
	(-.295)	(-.322)	(-.288)	(-.417)	(-.499)	(-.761)		
	(.941)	(.721)	(.675)	(1.232)	(.435)	(.171)		
Survey Weight	.126	-.014	.087	.024	-.090	-.264		
WB GDP Weight	.873	1.014	.912	.975	1.090	1.264		
1992-1997	.167	.074	.035	.034	.332	.062	165	88
	(-.092)	(-.165)	(-.282)	(-.280)	(-.168)	(-.335)		
	(.658)	(.527)	(.573)	(.522)	(1.676)	(.604)		
Survey Weight	.143	.069	.033	.033	.249	.059		
WB GDP Weight	.856	.930	.966	.966	.750	.940		
1998-2003	.003	-.112	-.011	-.086	.063	-.082	234	98
	(-.284)	(-.405)	(-.320)	(-.453)	(-.296)	(-.407)		
	(.344)	(.353)	(.365)	(.537)	(.477)	(.358)		
Survey Weight	.003	-.127	-.011	-.094	.059	-.090		
WB GDP Weight	.996	1.127	1.011	1.094	.940	1.090		
2004-2010	.464	.074	.469	.147	.349	.325	302	103
	(-.098)	(-.230)	(-.159)	(-.280)	(-.229)	(-.418)		
	(1.197)	(.686)	(1.291)	(1.375)	(1.698)	(3.758)		
Survey Weight	.317	.069	.319	.128	.258	.245		
WB GDP Weight	.682	.930	.680	.871	.741	.754		

Table III presents estimates and 95% confidence intervals for $\hat{\omega} = \hat{\gamma}_{surveys}/\hat{\gamma}_{NA}$ the ratio of the weight of log survey means per capita to the weight of log national accounts GDP per capita in the optimal unbiased proxy z_i for the mean of the true income distribution, as well as the values of the estimated weights under the assumption A5a. Each row corresponds to estimating the weight $\hat{\omega}$ for a different subsample of the baseline sample: either restricting to observations in a specific region or to observations in a specific year range. Data definitions, inference procedures and sample selection are as in Table I. The baseline specification corresponds to the specification in the bolded cell of Table II. Each column corresponds to a different specification in which either countries are weighted equally instead of surveys, or urban and rural control variables are included (log total population, log percentage rural population, log percentage urban population), or all control variables are included (log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP) or both.

Table IV

Developing World Poverty Estimates								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	1992	2005	2006	2007	2008	2009	2010	Ratio 2010-1992
Survey Weight = 1 (CR 2010)	.421	.258	.247	.237	.227	.214	.205	.487
GDP Weight = 1 (PSiM 2009)	.094	.050	.047	.043	.041	.039	.037	.400
Baseline	.118	.061	.057	.052	.049	.047	.045	.381
	(.087)	(.047)	(.044)	(.041)	(.039)	(.037)	(.036)	(.365)
	(.156)	(.079)	(.074)	(.068)	(.064)	(.060)	(.057)	(.409)
Scale Normalized to GDP	.099	.051	.048	.044	.042	.040	.038	.387
	(.092)	(.050)	(.046)	(.043)	(.040)	(.039)	(.037)	(.373)
	(.108)	(.054)	(.051)	(.047)	(.044)	(.042)	(.040)	(.405)
Scale Normalized to Surveys	.289	.170	.158	.146	.138	.130	.121	.420
	(.258)	(.154)	(.143)	(.132)	(.124)	(.117)	(.110)	(.407)
	(.332)	(.190)	(.177)	(.164)	(.155)	(.145)	(.136)	(.437)
Year-spec. Weights Recursive Scale	.119	.059	.057	.053	.049	.047	.042	.354
	(.107)	(.055)	(.052)	(.048)	(.047)	(.043)	(.039)	(.311)
	(.131)	(.065)	(.065)	(.060)	(.052)	(.053)	(.046)	(.404)
Region-spec Weights	.101	.056	.052	.048	.045	.043	.040	.460
	(.039)	(.031)	(.028)	(.026)	(.024)	(.023)	(.022)	(.274)
	(.246)	(.115)	(.106)	(.096)	(.089)	(.082)	(.075)	(.708)
Urban / Rural Covariates	.108	.057	.053	.049	.046	.044	.042	.391
	(.076)	(.043)	(.040)	(.037)	(.035)	(.034)	(.032)	(.367)
	(.148)	(.075)	(.070)	(.064)	(.061)	(.057)	(.054)	(.431)
All Covariates	.104	.055	.051	.047	.044	.043	.040	.398
	(.071)	(.041)	(.039)	(.036)	(.034)	(.033)	(.031)	(.366)
	(.152)	(.077)	(.072)	(.066)	(.062)	(.058)	(.055)	(.441)
Add NA Consumption	.123	.064	.060	.055	.052	.049	.047	.382
	(.098)	(.057)	(.053)	(.049)	(.046)	(.044)	(.042)	(.351)
	(.146)	(.072)	(.067)	(.061)	(.057)	(.054)	(.051)	(.421)
Inequality Upper Bd. Gini + 1.96 * Med.SD	.163	.089	.083	.076	.071	.067	.064	.391
	(.126)	(.068)	(.063)	(.058)	(.055)	(.052)	(.049)	(.388)
	(.206)	(.115)	(.107)	(.098)	(.093)	(.087)	(.082)	(.399)
Conservative Forecasts								.540
								(.568)

Each row of Table IV presents estimates and 90% confidence intervals (5% and 95% confidence bounds) for developing world poverty rates in selected years using the estimated proxies z_i as the means of the country income distributions. Data definitions, inference procedures and sample selection for the sample used to compute the weights on national accounts and survey means in the construction of z_i are as in Table I. Poverty estimates are constructed using these weights for the whole sample of country-years of all countries not including the OECD and countries with no household surveys, and all years in the time period 1992-2010. Poverty estimates are obtained as the fraction of the population below $1.25 a day, with the income distribution assumed to be lognormal with mean equal to z_i and variance implied by the Gini coefficient from the corresponding household survey. All estimates obtained as means of corresponding bootstrapped distributions; estimates ratios need not equal exactly to ratios of estimates because of Jensen's inequality. Row 1 presents estimates in which z_i is set to the survey mean (as in CR (2010)). Row 2 presents estimates in which z_i is set to national accounts GDP per capita (as in PSiM (2009)). Row 3 presents the baseline specification, where the weights corresponds to the specification in the bolded cells of Table II and Table III, and Assumption A5a is invoked to fix the overall scale of z_i. Row 4 presents the baseline specification with the scale based on Assumption A5b. Row 5 presents the baseline specification with the scale based on Assumption A5c. Row 6 presents the baseline specification from Row 3 but with additional control variables for the estimation of the weights. (log total population, log percentage rural population, log percentage urban population). Row 7 presents the same specification as Row 6 but with further control variables for the estimation of the weights (log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP). Row 8 presents the same specification as Row 3 but adds national accounts consumption per capita as an additional component of the proxy z_i, and replaces assumption A5a with an analogous assumption to A5b in which the parameters associated with national accounts consumption are normalized. Row 9 presents the same specification as Row 3 but assumes all survey Gini coefficients are 5.37 Gini points higher than they are recorded to be in the household surveys. The conservative ratio in the last cell of Row 9 assumes that the survey Gini coefficients in 1992 are as reported, but the survey Gini coefficients in 2010 are 5.37 points higher than reported. We also present the upper bound for this ratio.

Table V (V)

Regional Poverty Estimates								
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Dev. World	East Asia	South Asia	Lat. Am.	SSA	MENA	Fmr USSR
Survey Weight = 1 (CR 2010)	Poverty 1992	.421	.512	.545	.129	.585	.074	.084
	Poverty 2010	.205	.093	.321	.058	.474	.048	.071
	Ratio 2010/1992	.487	.182	.588	.455	.811	.651	.841
GDP Weight = 1 (PSiM 2009)	Poverty 1992	.094	.081	.072	.026	.346	.003	.030
	Poverty 2010	.037	.002	.008	.017	.217	.003	.009
	Ratio 2010/1992	.400	.031	.119	.673	.628	1.037	.327
Baseline	Poverty 1992	.118	.115	.105	.033	.374	.005	.031
	Poverty 2010	.045	.004	.016	.020	.244	.005	.015
	Ratio 2010/1992	.381	.040	.149	.625	.650	.962	.478
	Ratio 2010/1992 UB	(.409)	(.054)	(.201)	(.690)	(.681)	(1.060)	(.691)
Scale Normalized to GDP	Poverty 1992	.099	.092	.080	.026	.347	.003	.027
	Poverty 2010	.038	.002	.010	.016	.219	.003	.012
	Ratio 2010/1992	.387	.032	.131	.638	.631	1.072	.455
	Ratio 2010/1992 UB	(.405)	(.033)	(.152)	(.684)	(.639)	(1.123)	(.668)
Scale Normalized to Surveys	Poverty 1992	.289	.321	.334	.118	.554	.052	.075
	Poverty 2010	.121	.046	.107	.076	.435	.033	.043
	Ratio 2010/1992	.420	.146	.319	.643	.785	.640	.571
	Ratio 2010/1992 UB	(.437)	(.163)	(.353)	(.688)	(.795)	(.654)	(.639)
Year-spec. Weights Recursive Scale	Poverty 1992	.119	.105	.106	.032	.401	.005	.040
	Poverty 2010	.042	.005	.019	.017	.218	.004	.019
	Ratio 2010/1992	.354	.053	.182	.547	.545	.826	.509
	Ratio 2010/1992 UB	(.404)	(.077)	(.283)	(.667)	(.663)	(.999)	(.754)
Region-spec Weights	Poverty 1992	.101	.098	.093	.039	.299	.007	.031
	Poverty 2010	.040	.008	.028	.023	.179	.005	.013
	Ratio 2010/1992	.460	.039	.363	.613	.593	.902	.434
	Ratio 2010/1992 UB	(.708)	(.106)	(1.744)	(.772)	(.649)	(1.192)	(.755)
Urban / Rural Covariates	Poverty 1992	.108	.101	.092	.030	.362	.004	.030
	Poverty 2010	.042	.004	.013	.019	.233	.004	.012
	Ratio 2010/1992	.391	.036	.138	.647	.641	.995	.415
	Ratio 2010/1992 UB	(.431)	(.051)	(.190)	(.728)	(.675)	(1.105)	(.653)
All Covariates	Poverty 1992	.104	.096	.086	.029	.357	.004	.030
	Poverty 2010	.040	.003	.012	.018	.228	.004	.012
	Ratio 2010/1992	.398	.035	.135	.660	.637	1.010	.388
	Ratio 2010/1992 UB	(.441)	(.052)	(.195)	(.744)	(.678)	(1.123)	(.670)
Add NA Consumption	Poverty 1992	.123	.123	.109	.037	.378	.006	.033
	Poverty 2010	.047	.005	.017	.023	.250	.005	.016
	Ratio 2010/1992	.382	.047	.161	.619	.662	.924	.490
	Ratio 2010/1992 UB	(.421)	(.054)	(.190)	(.671)	(.696)	(1.047)	(.675)
Inequality Upper Bd. Gini + 1.96 * Med.SD	Poverty 1992	.163	.170	.161	.057	.422	.015	.043
	Poverty 2010	.064	.013	.037	.035	.296	.011	.022
	Ratio 2010/1992	.391	.078	.223	.617	.699	.748	.523
	Ratio 2010/1992 UB	(.399)	(.102)	(.293)	(.661)	(.728)	(.794)	(.660)
Conservative Forecasts	Extreme Ratio 2010/1992	.540	.116	.346	1.073	.789	2.181	.725
	Extreme Ratio 2010/1992 UB	(.568)	(.139)	(.417)	(1.172)	(.811)	(2.598)	(.957)

Each row of Table V presents estimates and 90% confidence intervals (5% and 95% confidence bounds) for poverty rates in selected developing world regions using the estimated proxies z_i as the means of the country income distributions. Data definitions, inference procedures and sample selection for the sample used to compute the weights on national accounts and survey means in the construction of z_i are as in Table I. Poverty estimates are constructed using these weights for the whole sample of country-years of all countries not including the OECD and countries with no household surveys, and all years in the time period 1992-2010. Poverty estimates are obtained as the fraction of the population below $1.25 a day, with the income distribution assumed to be lognormal with mean equal to z_i and variance implied by the Gini coefficient from the corresponding household survey. All estimates obtained as means of corresponding bootstrapped distributions; estimates ratios need not equal exactly to ratios of estimates because of Jensen's inequality. Row 1 presents estimates in which z_i is set to the survey mean (as in CR (2010)). Row 2 presents estimates in which z_i is set to national accounts GDP per capita (as in PSiM (2009)). Row 3 presents the baseline specification, where the weights corresponds to the specification in the bolded cells of Table II and Table III, and Assumption A5a is invoked to fix the overall scale of z_i. Row 4 presents the baseline specification with the scale based on Assumption A5b. Row 5 presents the baseline specification with the scale based on Assumption A5c. Row 6 presents the baseline specification from Row 3 but with additional control variables for the estimation of the weights. (log total population, log percentage rural population, log percentage urban population). Row 7 presents the same specification as Row 6 but with further control variables for the estimation of the weights (log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP). Row 8 presents the same specification as Row 3 but adds national accounts consumption per capita as an additional component of the proxy z_i, and replaces assumption A5a with an analogous assumption to A5b in which the parameters associated with national accounts consumption are normalized. Row 9 presents the same specification as Row 3 but assumes all survey Gini coefficients are 5.37 Gini points higher than they are recorded to be in the household surveys. The conservative ratio in the last cell of Row 9 assumes that the survey Gini coefficients in 1992 are as reported, but the survey Gini coefficients in 2010 are 5.37 points higher than reported. We also present the 95% upper bound for this ratio.

Table VI

Developing World True Income Estimates							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	1992	2005	2006	2007	2008	2009	2010
Survey Weight = 1	1149	1440	1526	1611	1681	1735	1794
GDP Weight = 1	2905	4286	4578	4916	5133	5199	5442
Baseline	2549	3701	3948	4228	4414	4479	4680
	(2122)	(3002)	(3197)	(3411)	(3559)	(3622)	(3775)
	(3034)	(4501)	(4809)	(5170)	(5398)	(5464)	(5724)
Scale Normalized	2851	4161	4447	4772	4988	5061	5295
to GDP	(2771)	(3984)	(4264)	(4572)	(4787)	(4871)	(5092)
	(2917)	(4315)	(4607)	(4950)	(5167)	(5230)	(5477)
Scale Normalized	1229	1691	1780	1881	1946	1976	2046
to Surveys	(1178)	(1624)	(1707)	(1802)	(1866)	(1896)	(1963)
	(1273)	(1757)	(1851)	(1960)	(2031)	(2061)	(2136)
Year-spec. Weights	2724	3725	3583	3675	4170	3994	4414
Recursive Scale	(2400)	(3590)	(3160)	(3344)	(3900)	(3642)	(4150)
	(3315)	(3910)	(4030)	(4028)	(4449)	(4288)	(4673)
Region-spec Weights	8783	13144	15732	19774	21687	22926	23305
	(2168)	(2965)	(3120)	(3319)	(3465)	(3481)	(3641)
	(13243)	(27447)	(31861)	(37952)	(40779)	(42570)	(44475)
Urban / Rural	2717	3978	4246	4555	4755	4821	5043
Covariates	(2186)	(3108)	(3310)	(3533)	(3687)	(3751)	(3911)
	(3344)	(5014)	(5363)	(5778)	(6035)	(6101)	(6400)
All Covariates	2818	4147	4428	4755	4965	5030	5265
	(2158)	(3062)	(3261)	(3480)	(3631)	(3694)	(3851)
	(3484)	(5248)	(5615)	(6056)	(6326)	(6392)	(6709)
Add NA	2361	3391	3608	3853	4015	4077	4252
Consumption	(2232)	(3164)	(3335)	(3567)	(3711)	(3781)	(3930)
	(2508)	(3641)	(3903)	(4160)	(4335)	(4387)	(4584)
Inequality Upper Bd.	2549	3701	3948	4228	4414	4479	4680
Gini + 1.96 * Med.SD	(2122)	(3002)	(3197)	(3411)	(3559)	(3622)	(3775)
	(3034)	(4501)	(4809)	(5170)	(5398)	(5464)	(5724)

Each row of Table VI presents estimates and 90% confidence intervals (5% and 95% confidence bounds) for developing world true income per capita (the population-weighted average of the z_i's) in selected years. Data definitions, inference procedures and sample selection for the sample used to compute the weights on national accounts and survey means in the construction of z_i are as in Table I. Row 1 presents estimates in which z_i is set to the survey mean (as in CR (2010)). Row 2 presents estimates in which z_i is set to national accounts GDP per capita (as in PSiM (2009)). Row 3 presents the baseline specification, where the weights corresponds to the specification in the bolded cells of Table II and Table III, and Assumption A5a is invoked to fix the overall scale of z_i. Row 4 presents the baseline specification with the scale based on Assumption A5b. Row 5 presents the baseline specification with the scale based on Assumption A5c. Row 6 presents the baseline specification with additional control variables for the estimation of the weights. (log total population, log percentage rural population, log percentage urban population). Row 7 presents further control variables for the estimation of the weights (log consumption share, log capital formation as percent of GDP, log shares of GDP in agriculture, manufacturing and services, log export share, log import share, and log government expenditure share of GDP). Row 8 presents the same specification as Row 3 but adds national accounts consumption per capita as an additional component of the proxy z_i, and replaces assumption A5a with an analogous assumption to A5b in which the parameters associated with national accounts consumption are normalized. Row 9 presents the same specification as Row 3 but assumes all survey Gini coefficients are 5.37 Gini points higher than they are recorded to be in the household surveys.

Table VII

(VII)

Regional GDP per Capita Estimates								
		(1)	(2)	(3)	(4)	(5)	(6)	(7)
		Dev. World	East Asia	South Asia	Lat. Am.	SSA	MENA	Fmr USSR
Survey Weight = 1 (CR 2010)	GDP per capita in 1992	1149	612	509	2653	664	1760	2985
	GDP per capita in 2005	1440	1281	662	3305	757	2015	2899
	GDP per capita in 2010	1794	1805	730	4228	810	2070	3993
GDP Weight = 1 (PSiM 2009)	GDP per capita in 1992	2905	1672	1250	7384	1547	5100	7624
	GDP per capita in 2005	4286	4020	2098	8817	1756	6704	8599
	GDP per capita in 2010	5442	6164	2810	10115	2017	6821	10322
Baseline	GDP per capita in 1992	2549	1460	1104	6424	1366	4420	6681
	GDP per capita in 2005	3701	3456	1800	7721	1554	5711	7426
	GDP per capita in 2010	4680	5243	2352	8968	1763	5816	9058
Scale Normalized to GDP	GDP per capita in 1992	2851	1591	1191	7305	1500	4976	7623
	GDP per capita in 2005	4161	3850	1965	8836	1713	6465	8491
	GDP per capita in 2010	5295	5909	2582	10325	1947	6591	10443
Scale Normalized to Surveys	GDP per capita in 1992	1229	824	669	2717	747	1996	2787
	GDP per capita in 2005	1691	1656	986	3153	829	2445	3015
	GDP per capita in 2010	2046	2308	1217	3562	920	2478	3543
Year-spec. Weights Recursive Scale	GDP per capita in 1992	2724	1572	1165	6996	1448	4852	7160
	GDP per capita in 2005	3725	3479	1812	7762	1562	5756	7475
	GDP per capita in 2010	4414	4798	2077	9078	1746	5364	8956
Region-spec Weights	GDP per capita in 1992	8783	7554	11691	6280	2430	21022	7211
	GDP per capita in 2005	13144	17095	9882	7534	2730	40275	8095
	GDP per capita in 2010	23305	33575	25960	8751	3378	41843	9763
Urban / Rural Covariates	GDP per capita in 1992	2717	1560	1173	6878	1452	4742	7126
	GDP per capita in 2005	3978	3723	1942	8237	1650	6184	7982
	GDP per capita in 2010	5043	5681	2572	9506	1884	6295	9654
All Covariates	GDP per capita in 1992	2818	1620	1215	7153	1504	4936	7395
	GDP per capita in 2005	4147	3886	2028	8547	1708	6475	8320
	GDP per capita in 2010	5265	5950	2709	9827	1959	6590	10012
Add NA Consumption	GDP per capita in 1992	2361	1397	1074	5812	1299	4047	6053
	GDP per capita in 2005	3391	3199	1713	6935	1470	5169	6645
	GDP per capita in 2010	4252	4767	2210	8016	1660	5267	8063
Inequality Upper Bd. Gini + 1.96 * Med.SD	GDP per capita in 1992	2549	1460	1104	6424	1366	4420	6681
	GDP per capita in 2005	3701	3456	1800	7721	1554	5711	7426
	GDP per capita in 2010	4680	5243	2352	8968	1763	5816	9058

Each row of Table VII presents estimates and 90% confidence intervals (5% and 95% confidence bounds) for true income per capita (the population-weighted average of the z_i's) in selected developing world regions. Data definitions, inference procedures and sample selection for the sample used to compute the weights on national accounts and survey means in the construction of z_i are as in Table I. All row specifications as in Table VI.

9 Figures

Figure I (I)

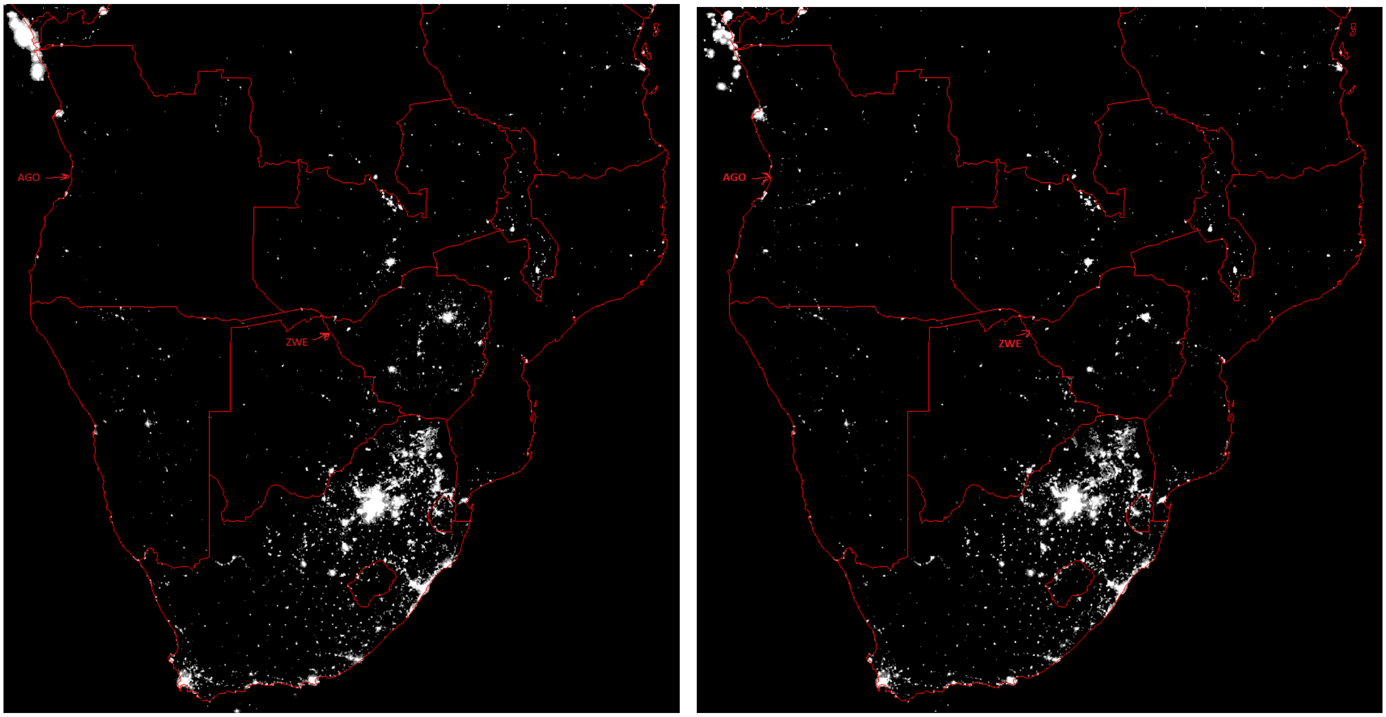

Southern Africa, 2000 Southern Africa, 2009

Data Source: NOAA. The symbols "AGO", "ZWE" and "BWA" show Angola, Zimbabwe and Botswana respectively (the Zimbabwe symbol placed in Botswana near its Zimbabwean border to avoid masking Zimbabwean lights).

Figure II (II)

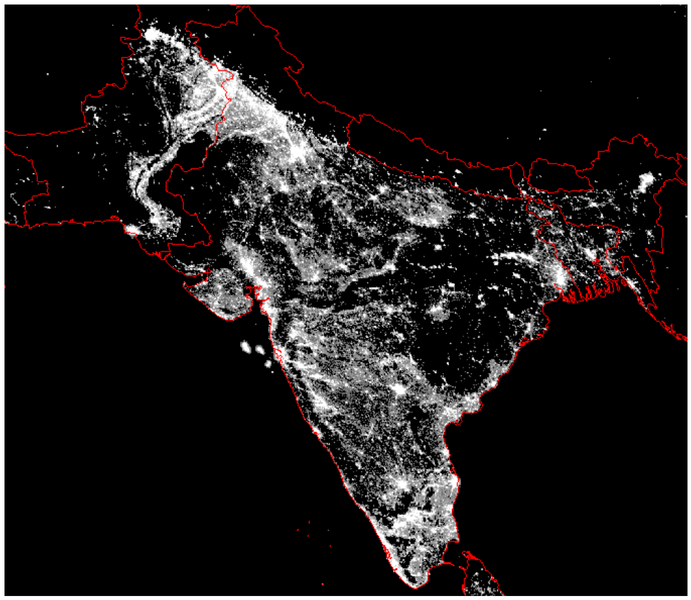

India, 1994

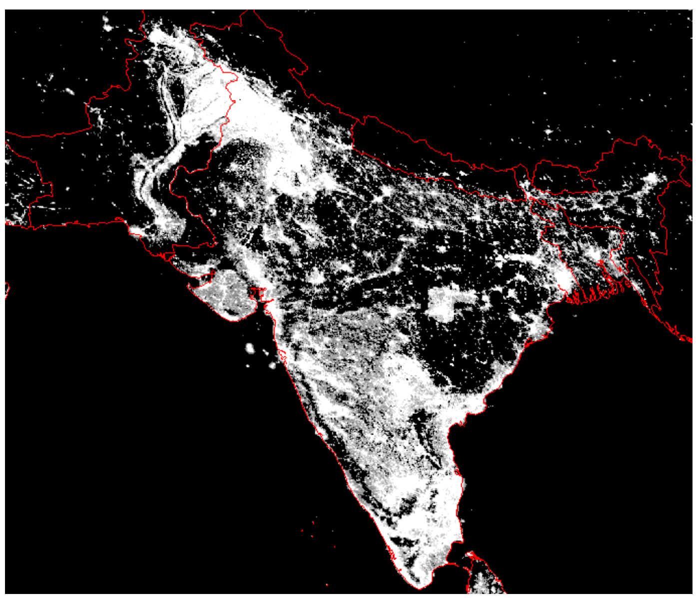

India, 2010

Data Source: NOAA.

Figure III (III)

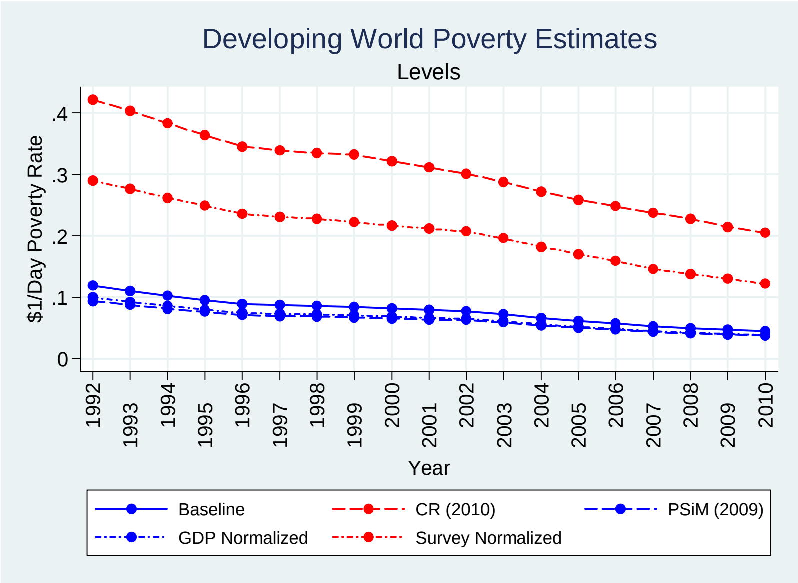

Note: See Table IV for data and series descriptions.

Figure IV (IV)

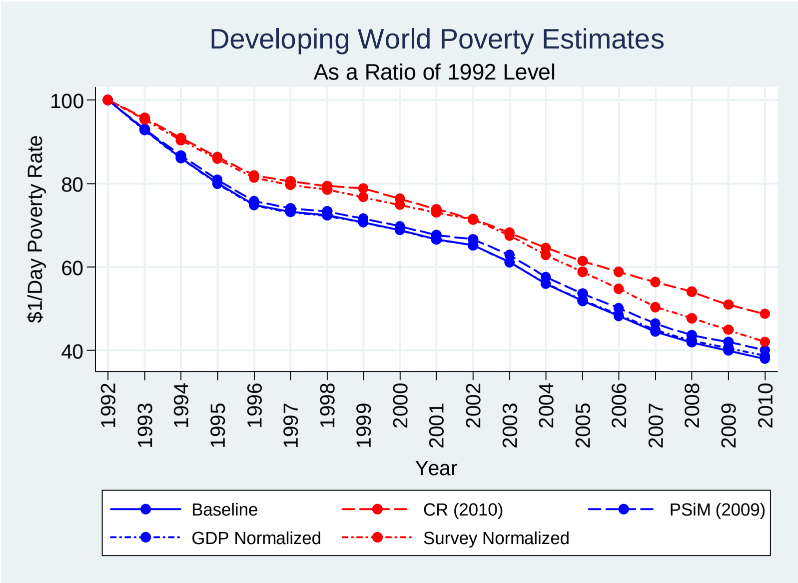

Note: See Table IV for data and series descriptions.

10 Appendix Tables

Table AI (AI)

Summary Statistics				
Series	Mean Whole World	SD Whole World	Mean Base Sample	SD Base Sample
Log Lights per Capita	18.06	1.66	18.11	1.21
Log WB GDP per Capita, PPP	8.56	1.28	8.41	.88
Log Survey Mean, PPP	7.54	.74	7.54	.74
Log WB NA Consumption per Capita, PPP	8.33	1.19	8.22	.83
Log Fraction Rural Population	3.64	.72	3.71	.51
Log Total Population	15.13	2.31	16.30	1.53
Log Fraction Urban Population	3.87	.54	3.90	.45
Log Services Share of GDP	3.95	.33	3.95	.23
Log Agricultural Share of GDP	2.26	1.17	2.49	.70
Log Export Share of GDP	3.47	.67	3.47	.54
Log Import Share of GDP	3.69	.56	3.66	.52
Log Manufacturing Share of GDP	2.47	.66	2.75	.44
Log Consumption Share of GDP	4.16	.29	4.21	.20
Log Government Expenditure Share of GDP	2.69	.41	2.60	.36
Log Gross Capital Formation Share of GDP	3.05	.40	3.08	.31

Note: Table AI presents summary statistics of key variables in the analysis. "Whole World" refers to all countries and years in the universe of countries and from 1992 to 2010. "Base Sample" refers to the sample of 701 country-years for which both lights data and survey means are availaible and which is used to estimated optimal weights. Data on lights from the NOAA. All other data from the World Bank's World Development Indicators.

Table AII (AII)

Countries included in Calibration Sample									
Country	No. Surv.	First Year	Last Year	Log GDP First Yr.	Log GDP Last Yr.	Log Lights First Yr.	Log Lights Last Yr.	Log Surv. First Yr.	Log Surv. Last Yr.
Albania	5	1997	2008	8.18	8.88	17.19	18.37	7.50	7.64
Algeria	1	1995	1995	8.63	8.63	19.00	19.00	7.27	7.27
Angola	2	2000	2009	7.81	8.54	16.70	17.41	6.62	6.57
Argentina	19	1992	2010	9.12	9.57	18.88	19.73	8.48	8.99
Armenia	11	1996	2010	7.51	8.49	17.21	18.59	7.16	7.18
Azerbaijan	3	1995	2008	7.52	8.99	18.43	18.52	6.95	7.78
Bangladesh	5	1992	2010	6.65	7.30	15.79	16.17	6.02	6.42
Belarus	13	1993	2010	8.57	9.43	18.59	19.96	7.80	8.70
Belize	7	1993	1999	8.53	8.53	18.90	19.07	7.97	7.73
Benin	1	2003	2003	7.21	7.21	16.09	16.09	6.45	6.45
Bhutan	2	2003	2007	8.08	8.34	16.56	17.08	7.04	7.21
Bolivia	10	1993	2008	8.07	8.33	18.39	18.37	7.79	7.85
Bosnia Herzegovina	3	2001	2007	8.56	8.88	18.95	18.82	8.34	8.64
Botswana	1	1994	1994	8.88	8.88	18.10	18.10	7.33	7.33
Brazil	16	1992	2009	8.85	9.15	18.56	18.94	7.65	8.38
Bulgaria	7	1992	2007	8.78	9.32	18.56	18.82	8.57	8.09
Burkina Faso	4	1994	2009	6.53	6.98	15.58	16.04	6.19	6.51
Burundi	3	1992	2006	6.55	6.20	14.77	14.62	5.74	5.85
Cambodia	5	1994	2009	6.66	7.53	14.44	15.69	6.51	6.87
Cameroon	3	1996	2007	7.43	7.61	16.28	15.96	6.88	7.23
Cape Verde	1	2002	2002	7.76	7.76	17.78	17.78	7.28	7.28
Cent. African Rep.	3	1992	2008	6.62	6.55	15.68	14.71	5.69	6.42
Chad	1	2003	2003	6.84	6.84	14.88	14.88	6.20	6.20
Chile	8	1992	2009	8.99	9.53	18.26	18.88	8.22	8.68
China	7	1993	2009	7.31	8.73	16.96	17.88	6.34	7.47
Colombia	14	1992	2010	8.74	9.04	18.26	18.84	7.96	8.12
Comoros	1	2004	2004	6.94	6.94	15.10	15.10	7.03	7.03
Congo	1	2005	2005	8.12	8.12	17.93	17.93	6.47	6.47
Congo, DRC	1	2006	2006	5.64	5.64	15.39	15.39	5.56	5.56
Costa Rica	18	1992	2009	8.80	9.22	18.69	18.92	7.82	8.49
Cote d'Ivoire	5	1993	2008	7.45	7.41	16.68	17.21	6.94	6.95

Note: Table AII presents a list of countries and relevant statistics for the calibration sample of country-years, based on which we calculate weights on national accounts and survey means in the optimal proxy for log true income per capita. We present the number of surveys each country has in the sample, the years of the earliest and latest survey, and values of log World Bank GDP per capita (PPP-adjusted), log survey mean, and log lights per capita (NOAA) corresponding to these years.

Table AII (cont.)

Countries included in Calibration Sample									
Country	No. Surv.	First Year	Last Year	Log GDP First Yr.	Log GDP Last Yr.	Log Lights First Yr.	Log Lights Last Yr.	Log Surv. First Yr.	Log Surv. Last Yr.
Croatia	6	1998	2008	9.37	9.75	19.45	19.80	8.72	9.12
Czech Republic	2	1993	1996	9.57	9.70	19.28	19.70	8.54	8.69
Djibouti	1	2002	2002	7.47	7.47	16.31	16.31	7.02	7.02
Dominican Republic	14	1992	2010	8.32	9.03	17.76	18.44	7.91	8.03
Ecuador	11	1994	2010	8.62	8.88	18.56	19.30	7.68	8.07
Egypt	4	1996	2008	8.19	8.55	18.68	18.98	7.06	7.22
El Salvador	13	1995	2009	8.43	8.68	18.00	18.14	7.75	7.81
Estonia	8	1993	2004	8.90	9.62	19.42	19.59	8.09	8.21
Ethiopia	3	1995	2005	6.18	6.45	14.74	14.82	6.29	6.42
Fiji	2	2003	2009	8.31	8.34	17.57	17.54	6.99	7.43
Gabon	1	2005	2005	9.47	9.47	19.24	19.24	7.49	7.49
Georgia	14	1996	2010	7.60	8.42	16.92	18.63	7.59	7.16
Ghana	3	1992	2006	6.84	7.13	16.91	16.96	6.37	6.87
Guatemala	6	1998	2006	8.25	8.33	17.78	17.68	7.65	7.78
Guinea	3	1994	2007	6.71	6.88	15.73	15.34	6.21	6.52
Guinea-Bissau	2	1993	2002	7.11	6.93	15.61	14.73	6.51	6.36
Guyana	2	1993	1998	7.56	7.80	17.65	18.18	7.82	7.67
Haiti	1	2001	2001	7.00	7.00	15.40	15.40	6.50	6.50
Honduras	17	1992	2009	7.91	8.15	17.46	18.17	7.15	7.79
Hungary	8	1993	2007	9.31	9.78	18.83	18.98	8.33	8.47
India	3	1994	2010	7.19	8.01	17.11	17.86	6.32	6.58
Indonesia	7	1993	2010	7.78	8.26	17.02	17.65	6.26	6.90
Iran	3	1994	2005	8.80	9.13	19.12	19.27	7.93	7.77
Iraq	1	2007	2007	8.01	8.01	18.65	18.65	7.17	7.17
Jamaica	6	1993	2004	8.90	8.85	18.72	18.76	7.31	8.11
Jordan	6	1992	2010	8.12	8.56	18.92	19.71	7.64	7.90
Kazakhstan	10	1993	2009	8.59	9.24	19.58	19.50	7.33	7.76
Kenya	4	1992	2005	7.19	7.20	16.07	15.73	7.01	6.66
Kyrgyzstan	10	1993	2010	7.40	7.61	18.55	18.82	7.63	7.30
Laos	4	1992	2008	6.89	7.62	15.64	16.89	6.25	6.62
Latvia	11	1993	2009	8.68	9.46	18.51	18.96	7.76	8.47

Note: Table AII presents a list of countries and relevant statistics for the calibration sample of country-years, based on which we calculate weights on national accounts and survey means in the optimal proxy for log true income per capita. We present the number of surveys each country has in the sample, the years of the earliest and latest survey, and values of log World Bank GDP per capita (PPP-adjusted), log survey mean, and log lights per capita (NOAA) corresponding to these years.

Table AII (cont.)

Country	No. Surv.	First Year	Last Year	Log GDP First Yr.	Log GDP Last Yr.	Log Lights First Yr.	Log Lights Last Yr.	Log Surv. First Yr.	Log Surv. Last Yr.
Countries included in Calibration Sample									
Lesotho	3	1993	2003	6.86	7.06	16.50	16.66	6.61	6.76
Liberia	1	2007	2007	5.99	5.99	15.26	15.26	5.78	5.78
Lithuania	8	1993	2008	8.96	9.77	18.52	19.14	7.31	8.58
Macedonia	10	1998	2010	8.82	9.12	18.83	19.31	7.74	8.04
Madagascar	6	1993	2010	6.82	6.76	15.20	15.53	6.07	5.81
Malawi	3	1998	2010	6.51	6.65	16.25	16.64	5.86	6.27
Malaysia	6	1992	2009	8.95	9.47	18.12	19.04	8.01	8.47
Maldives	2	1998	2004	8.31	8.67	14.13	14.22	7.80	7.65
Mali	4	1994	2010	6.48	6.87	15.53	16.67	5.65	6.32
Mauritania	5	1993	2008	7.48	7.70	16.72	17.05	6.74	6.92
Mexico	11	1992	2010	9.24	9.43	18.81	19.34	8.12	8.14
Moldova	14	1992	2010	7.90	7.93	19.14	18.72	6.94	7.71
Mongolia	4	1995	2008	7.60	8.17	17.70	18.02	6.87	7.49
Montenegro	6	2005	2010	9.01	9.22	19.00	19.72	8.08	8.24
Morocco	3	1999	2007	7.97	8.24	17.71	17.99	7.35	7.56
Mozambique	3	1996	2008	6.03	6.63	15.93	16.35	5.88	6.32
Namibia	2	1993	2004	8.32	8.55	18.31	18.33	7.47	7.46
Nepal	3	1996	2010	6.72	6.98	15.66	15.89	6.11	6.70
Nicaragua	4	1993	2005	7.72	8.01	17.49	17.53	7.16	7.50
Niger	4	1992	2008	6.44	6.48	15.70	15.54	6.02	6.45
Nigeria	4	1992	2010	7.28	7.66	17.98	17.51	6.17	6.17
Pakistan	6	1997	2008	7.49	7.74	17.74	17.70	6.32	6.67
Panama	11	1995	2010	8.87	9.44	18.63	19.06	8.09	8.16
Papua New Guinea	1	1996	1996	7.76	7.76	17.15	17.15	6.94	6.94
Paraguay	13	1995	2010	8.38	8.43	18.76	19.10	8.13	8.14
Peru	15	1994	2010	8.51	9.05	17.80	18.51	7.41	8.06
Philippines	6	1994	2009	7.81	8.12	16.58	16.66	6.90	7.12
Poland	15	1992	2010	8.95	9.76	18.77	20.47	8.08	8.42
Romania	13	1992	2010	8.75	9.29	17.64	19.38	7.92	7.87
Russia	12	1993	2009	9.14	9.51	19.78	19.81	8.19	8.58
Rwanda	2	2000	2006	6.48	6.79	14.96	14.53	6.13	6.22

Note: Table AII presents a list of countries and relevant statistics for the calibration sample of country-years, based on which we calculate weights on national accounts and survey means in the optimal proxy for log true income per capita. We present the number of surveys each country has in the sample, the years of the earliest and latest survey, and values of log World Bank GDP per capita (PPP-adjusted), log survey mean, and log lights per capita (NOAA) corresponding to these years.

Table AII (cont.)

Country	No. Surv.	First Year	Last Year	Log GDP First Yr.	Log GDP Last Yr.	Log Lights First Yr.	Log Lights Last Yr.	Log Surv. First Yr.	Log Surv. Last Yr.
Countries included in Calibration Sample									
Senegal	3	1994	2005	7.23	7.42	16.43	16.50	6.39	6.68
Serbia	9	2002	2010	8.87	9.16	19.24	20.02	8.30	8.19
Seychelles	2	2000	2007	9.84	9.95	18.75	18.55	8.63	8.62
Sierra Leone	1	2003	2003	6.41	6.41	14.17	14.17	6.42	6.42
Slovakia	8	1992	2009	9.22	9.87	19.53	19.12	8.48	8.39
Slovenia	5	1993	2004	9.59	10.02	19.02	19.23	8.80	9.01
South Africa	5	1993	2009	8.90	9.14	18.92	18.91	7.63	8.03
Sri Lanka	4	1996	2010	7.84	8.43	17.21	18.16	6.90	7.25
St. Lucia	1	1995	1995	9.01	9.01	19.00	19.00	7.07	7.07
Sudan	1	2009	2009	7.58	7.58	17.22	17.22	6.88	6.88
Suriname	1	1999	1999	8.52	8.52	18.98	18.98	7.71	7.71
Swaziland	3	1995	2010	8.30	8.58	18.07	18.76	6.02	6.86
Syria	1	2004	2004	8.29	8.29	18.81	18.81	7.39	7.39
Tajikistan	5	1999	2009	6.80	7.52	17.87	16.93	6.25	7.08
Tanzania	3	1992	2007	6.71	7.05	15.51	15.48	5.98	6.09
Thailand	10	1992	2010	8.41	8.94	17.67	19.07	7.45	7.88
The Gambia	2	1998	2003	7.29	7.35	16.20	15.72	6.22	6.89
Togo	1	2006	2006	6.77	6.77	16.00	16.00	6.51	6.51
Trinidad Tobago	1	1992	1992	9.28	9.28	19.18	19.18	7.71	7.71
Tunisia	4	1995	2010	8.50	9.04	18.80	19.37	7.52	7.92
Turkey	10	1994	2010	9.01	9.43	18.24	19.06	7.80	8.14
Turkmenistan	2	1993	1998	8.43	8.09	18.92	19.20	6.11	6.90
Uganda	6	1992	2009	6.35	7.02	14.83	15.16	6.11	6.70
Ukraine	13	1992	2010	8.80	8.70	19.14	19.26	7.97	8.25
Uruguay	5	2006	2010	9.21	9.44	18.82	19.55	8.38	8.61
Uzbekistan	3	1998	2003	7.34	7.48	18.76	18.42	6.82	6.42
Venezuela	10	1992	2006	9.27	9.27	19.32	19.23	7.88	7.87
Vietnam	6	1993	2008	6.97	7.86	15.79	17.39	6.17	6.93
Yemen	2	1998	2005	7.62	7.71	17.34	17.56	6.98	6.91
Zambia	7	1993	2010	7.10	7.24	17.32	17.80	6.22	6.14

Note: Table AII presents a list of countries and relevant statistics for the calibration sample of country-years, based on which we calculate weights on national accounts and survey means in the optimal proxy for log true income per capita. We present the number of surveys each country has in the sample, the years of the earliest and latest survey, and values of log World Bank GDP per capita (PPP-adjusted), log survey mean, and log lights per capita (NOAA) corresponding to these years.

Table AIII (AIII)

Comparison of PovcalNet and LIS						
Country	Year	PovcalNet (CR 2010) Inc. Concept	PovcalNet (CR 2010) Survey Mean	LIS (2013) Disposable Income Mean	NA (WB 2013) Consumption	NA (WB 2013) GDP
Brazil	2006	Income	3893.87	6043.15	7032.49	8753.23
Colombia	2004	Income	2091	3749.59	5871.28	7083.99
Estonia	2000	Consumption	3292.19	7381.67	8659.44	11512.50
Guatemala	2006	Income	2399.39	4884.08	4014.18	4175.75
Hungary	1999	Consumption	3376.67	8278	10000.29	13085.22
Mexico	1998	Income	2747.03	5549.40	8579.94	11030.44
Mexico	2000	Income	3330.11	6609.20	9260.07	11852.71
Mexico	2002	Income	3311.28	6634.41	9431.04	11621.00
Mexico	2004	Income	3498.71	6893.38	9221.80	11959.30
Peru	2004	Income	2553.59	4149.02	4742.52	6048.31
Poland	1999	Consumption	3580.32	8249.87	9043.51	11212.92
Poland	2004	Consumption	4087.91	8403.37	10941.17	13297.13

Note: Table AIII presents a list of survey means from PovcalNet (CR 2010) and from the Luxembourg Income Study (LIS 2013), as well as a list of national accounts consumption and GDP per capita from the World Bank for 12 country-years for which both a PovcalNet survey and a LIS survey is available. It also presents the corresponding income concept for the PovcalNet survey (the LIS income concept is household disposable income).