

A BAYESIAN MEASURE OF POVERTY IN THE DEVELOPING WORLD

BY ZHOU XUN

School of Economics, Jiangxi University of Finance and Economics

AND

MICHEL LUBRANO*

Aix-Marseille University, (Aix-Marseille School of Economics), GREQAM-CNRS & EHESS

We propose a new methodology to revise the international poverty line (IPL) after Ravallion *et al.* (2009) using the same database, but augmented with new variables to take into account social inclusion in the definition of poverty along the lines of Atkinson and Bourguignon (2001). We provide an estimation of the world income distribution and of the corresponding number of poor people in the developing world. Our revised IPL is based on an augmented two-regime model estimated using a Bayesian approach, which allows us to take into account uncertainty when defining the reference group of countries where the IPL applies. The influence of weighting by population is discussed, as well as the IPL revision proposed in Deaton (2010). We also discuss the impact of using the new 2011 PPP and the recent IPL revision made by the World Bank.

JEL Codes: C11, C21, I32

Keywords: Bayesian inference, poverty line, social inclusion, WDI

1. INTRODUCTION

The extent to which economic growth reduces poverty is widely debated. As Deaton (2005) says (see also references therein), the answer to this question basically depends on two other questions: Is the income of the poor growing faster than average income? And how is the poverty line fixed? If all the poor are very near the poverty line and if the latter is absolute, then economic growth will reduce poverty quite quickly. If these two conditions are not met, the result will be much more ambiguous and will depend on where in the income distribution

Note: The first idea for this paper grew out of participating in the conference “Pauvreté et décisions individuelles dans les pays en développement: que nous apprennent les études microéconométriques?” organized in Marseille on October 13, 2011. We are grateful to Christophe Muller for useful comments and discussions. The paper owes much to comments and suggestions made by Tony Atkinson on reading a previous version, which led to significant revisions. In memory of the kind person that Tony Atkinson was. The revised version of the paper was presented during the international workshop *Poverty and Well-Being*, held at Core, UCL, on October 26, 2016. Remarks and comments by the participants, two anonymous referees, the editor, and Denis de Crombrughe helped to improve the quality of our final product. We offer particular thanks to Marjorie Sweetko for the English revision. Of course, any remaining errors are solely ours. This work was carried out under the support of the A*MIDEX project (No. ANR-11-IDEX-0001-02) funded by the “Investissements d’Avenir” French government program, managed by the French National Research Agency (ANR).

*Correspondence to: Michel Lubrano, GREQAM, Centre de La Vieille Charité, 2 rue de la Charité, 13002 Marseille, France (michel.lubrano@univ-amu.fr).

individual incomes grow fastest. This ambiguity is illustrated, for instance, by the Kuznets curve. Inequality increases at the beginning of economic development (because the poorest individuals do not benefit much from economic growth) while, according to this model, inequality decreases in later phases of development, so that it becomes much easier to fight poverty. Of course, the speed at which poverty is reduced depends on where the poverty line is fixed, how it is determined, and how it has evolved. In order to study the evolution of poverty in the world, Atkinson and Bourguignon (2001) provide a useful discussion on the rationale behind the definition of a world poverty line.

A poverty line is an amount of income or a consumption level below which an individual or a representative household is declared to be poor. This is an important device, serving as the basis for many targeted economic policies such as the United Nations project *Objectives for the Millennium*, whose aim was to cut extreme poverty by half for 2015. There are several ways of defining a poverty line. An absolute poverty line is determined by reference to the cost of a given basket of goods. It is totally inelastic with respect to the country's mean income. Being poor in this context is equivalent to a lack of command over basic economic resources. A relative poverty line has a different purpose. Its aim is to situate a household within the income distribution, because it is defined in terms of a given percentage of a country's mean or median income. Its elasticity with respect to income is one. This type of poverty line is related to the concept of social inclusion as advanced in Atkinson and Bourguignon (2001). Being below a relative poverty line means being prevented from participating in ordinary, accepted social activities. The main example is inclusion in the functioning of the labor market. A subjective poverty line represents an intermediate case between these two polar cases, with intermediate income elasticity. It is obtained by processing opinion surveys containing either *the financial ease question* or *the minimum income question*, and has been used both for developed countries (see, e.g., Kapteyn *et al.*, 1988; Flik and van Praag, 1991), and for less developed countries (see, e.g., Pradhan and Ravallion, 2000). This poverty line represents a subjective aggregation of the different dimensions of a capability approach to poverty, which could be summarized as basic needs on one side and social inclusion for a given country on the other side (following Atkinson and Bourguignon, 2001).

The aim of this paper is to propose a Bayesian measure of poverty in the world—at least the main part of the world (in fact, the database we use represents more than 70 percent of the world population around 2001, naturally focusing on the less developed part of it). Basically, we need two tools: first, a correctly interpreted international poverty line and, second, an evaluation of the world income distribution. If we are concerned only by the less developed world, we need to consider an absolute poverty line. The World Bank provided the famous one-dollar-a-day international poverty line (IPL) in 1990 and Ravallion *et al.* (2009) proposed a method for reevaluating it, a method that we shall discuss in depth in this paper, as well as the last proposal made by the World Bank in 2015. If we want to evaluate poverty in the more developed world, the one-dollar approach is clearly inappropriate. A relative poverty line is needed, because roughly all the households are above the one-dollar-a-day poverty line. However, if we decided to apply a relative poverty line to the less developed world, we would obtain a

poverty line lower than the famous one-dollar-a-day threshold. So the decision on where to draw the two types of poverty line is important. Ravallion *et al.* (1991) and later Ravallion *et al.* (2009) produce an interesting graph which relates national poverty lines to national average income. Below a certain level of mean income, national poverty lines are constant around one dollar a day; above that threshold, national poverty lines tend to be a linear function of mean income. This graph suggests a two-regime econometric model which should help, first, to determine the allocation of countries between the two groups and, second, to determine the exact shape of the relation between national poverty lines and mean income. With this model, therefore, we might be able to answer some of the questions raised in Atkinson and Bourguignon (2001) concerning the exact proportion between basic needs and inclusion that a world poverty measurement should reflect. A Bayesian approach is fully justified to study our two-regime model. We know from Hansen (2000) that the classical approach leads to difficulties when deriving the asymptotic distribution of the threshold parameter, leading to confidence intervals that are too conservative. Moreover, the absence of pivotal statistics precludes the design of a simple bootstrap. With a classical approach as in Ravallion *et al.* (2009), the allocation between the two groups of countries is made conditional on the point estimate of the threshold parameter. With a Bayesian approach, we are able to produce a stochastic allocation so as to take better account of uncertainty. Moreover, we are able to produce a posterior density for the international poverty line (together with its posterior confidence interval).

The graphs produced by Ravallion *et al.* (1991) and Ravallion *et al.* (2009) are made possible by using PPP to obtain comparable figures for the different national poverty lines. Poverty lines are fixed in each country by official statistical agencies using national currencies. Using 2005 PPP, and more recently 2011 PPP, the World Bank converted these figures into PPP dollars. If the method of conversion were appropriate, an absolute international poverty line would reflect the minimum consumption cost covering basic human needs and activities in all the countries belonging to the least developed group. Thus the World Bank method, which consists of taking the mean of national poverty lines in this group, should be a valid way to define a universal absolute poverty line for the developing world. However, as discussed in Deaton (2010), the 2005 PPP revision changed both the geography and the extent of world poverty. The last 2011 revision of PPP led to even larger changes, so large that Ferreira *et al.* (2015) suggest keeping the 2005 PPP for some countries where the changes are too important and rather counterintuitive. So there is doubt about whether the 2005 PPP (and even more its 2011 revision) equalizes price levels across less developed countries. This raises the question of how national poverty lines vary across less developed countries, and why. And there might, of course, be reasons other than just the inappropriate nature of PPP (see, in particular, the discussion in Atkinson and Bourguignon, 2001).

The second tool we need is an estimation of the world income distribution. Sala-i-Martin (2006) has provided a non-parametric method which combines macro data and survey data. We do not intend to take this route, for two reasons. First, macro and survey data are usually not comparable. Second, survey data are available only for a limited number of countries. Further, non-parametric methods do not deal well with the tails of a distribution, which are of prime

importance for poverty analysis. We thus prefer to use a parametric approach, assuming that the income distribution of a single country can be described by a lognormal distribution. The world income distribution will then be formed by a mixture of lognormal distributions, the weight of each country being given by its population share. This idea is not new. We find it first in Atkinson and Bourguignon (2001), who calibrate the location of each lognormal density of PPP-corrected GDP per capita and make alternative assumptions on the dispersion. Holzmann *et al.* (2007) improved on this idea, using published Gini coefficients to determine the dispersion parameter.

The main original contributions of the paper are as follows. We first propose a revision of the international poverty line, no longer based on a simple average of national poverty lines for a given group. In a single statistical model, we both determine the composition of a group of less developed countries out of a much larger list of 74 countries, and estimate a common poverty line for these countries. The estimated common poverty line is not simply an average of national poverty lines, but at least partly takes into account the economic situation of each country. Using the same data as Ravallion *et al.* (2009), we obtain and justify a different specification which reflects the different concerns potentially raised by a poverty line. We conduct a detailed discussion showing why this database has to be used with the 2005 PPP and why converting it to the 2011 PPP would lead to data inconsistency. We also provide a discussion of the importance of population weighting and its influence on the final result. The Bayesian approach enables us to provide small sample confidence intervals and a posterior density for the revised IPL. In a second step, we propose an evaluation of the world income distribution based on a framework that is increasingly favored in the contemporary literature (see, e.g., Pinkovskiy and Sala-i-Martin, 2009). We contribute to this literature by providing a method allowing various data provided by the World Bank to be combined: different headcount rates (corresponding to different poverty lines), the Gini coefficient, and GDP per capita. Considering this distribution as given, we provide a Bayesian evaluation of the number of poor people in the world, using two different poverty lines: one for the less developed group, using our revised poverty line, and national relative poverty lines for the countries belonging to the other group. In so doing, we do not follow Sala-i-Martin (2006), who uses a single poverty line for all the countries, thus only dealing with absolute poverty. We are more in accordance with Atkinson and Bourguignon (2001), as we provide a measure of world poverty that combines both aspects of poverty: absolute poverty and relative poverty accounting for social inclusion.¹ The allocation of a country to one group or the other and the use of the corresponding poverty line are determined jointly inside the same Bayesian calculation. We are thus able to provide a posterior density for the number of poor people in the 74 countries of our sample. Above, we gave strong arguments supporting a Bayesian approach to this model. Finally, we provide a discussion about the recent revision by the World Bank of the IPL from \$1.25 using the 2005 PPP to \$1.90 using the 2011 PPP.

The paper is organized as follows. In Section 2, we discuss the model of Ravallion *et al.* (2009), propose a new formulation, show that there is too much

¹This reference and its relation to our work was kindly provided by A. B. Atkinson.

variability in national poverty lines, and argue that the new 2011 PPP introduces even more variability, shedding doubt on the validity of the new \$1.90 IPL recently released by the World Bank. We finally discuss the importance of weighting. Section 3 is devoted to Bayesian inference in a two-regime model and to the derivation of the posterior density of the IPL. We present our first empirical results in Section 4. In Section 5, we propose a new method of estimating the world income (consumption) distribution and discuss the method of deriving an alternative IPL suggested by Deaton (2010). We compare the IPL computed by Deaton (2010) to our revised IPL. In Section 6, we propose an evaluation of the number of poor people in the developing world, reflecting both a lack of command on basic needs and a lack of social inclusion. Section 7 concludes.

2. A REVISED COMMON POVERTY LINE FOR LESS DEVELOPED COUNTRIES

2.1. *What a Common Poverty Line Should Represent*

A national poverty line in less developed countries is usually defined as an “absolute poverty line” that focuses only on how much humans need to live, regardless of the national income distribution (see, e.g., Atkinson and Bourguignon, 2001). However, the minimum basket of goods ensuring a given level of physical and mental well-being varies from country to country, simply because living standards, traditions, habits, and other social characteristics are different (knowing that the PPP does not perfectly equalize the basic human needs, among the less developed countries). Can we explain these differences by observable characteristics, or are they just random?

In richer countries, once the basic needs are satisfied, individuals tend to desire a more expensive basket of goods—for example, more varied diets, suitable clothes, comfortable shelter, better health, and higher education—just to be like others and to be able to maintain a decent way of living (see, e.g., Atkinson, 1983, Ch. 10; Atkinson and Bourguignon, 2001). The definition of “poverty” in this case becomes more complex and is influenced largely by the perception of “economic inequality.” An individual who considers himself poor may not be facing a problem of survival, but may be suffering either from an envy-based comparison with what others in his surroundings possess or from a lack of social inclusion. The latter definition of poverty line is called a “relative poverty line” and corresponds to a position in the income distribution.

Where can the limit between these two definitions of a poverty line be set? Which countries are considered as being sufficiently rich to afford a relative poverty line, and which are the others? Ravallion *et al.* (1991) showed that official national poverty lines vary little in comparison with mean consumption per capita for less developed countries, while above a critical level of mean consumption per capita, national official poverty lines have a much stronger elasticity with respect to consumption. Based on that previous finding, Ravallion and Chen (2001, 2004) proposed an IPL (a worldwide absolute poverty line) of “\$1.00 per day” (\$1.08 at 1993 PPP).

In a more recent paper, Ravallion *et al.* (2009) clearly identify two groups of countries in a new dataset covering 74 developing countries, with data collected

over the period from 1988 to 2005. They estimate a non-linear regression relating national official poverty lines z_i to national mean consumption per capita C_i , imposing a zero consumption coefficient for the group of less developed countries, and thus leading implicitly to an absolute definition of the corresponding poverty line. Their model is equivalent to the following:

$$(1) \quad z_i = s_i(\alpha_1 + \gamma_1 C_i) + (1 - s_i)(\alpha_2 + \gamma_2 C_i) + \epsilon_i,$$

where s_i is equal to an indicator function $\mathbb{1}(C_i < \theta)$, which is one for countries below a mean consumption of θ and zero otherwise. For less developed countries, the elasticity of the poverty line with respect to mean consumption is assumed to be zero, that is, $\gamma_1 = 0$. In this case, α_1 corresponds to the mean of the dependent variable when $s_i = 1$ and is taken as an estimate of the revised IPL. Using this model with a fixed $\theta = \$60$, the revised IPL rises to \$1.25 per day at 2005 PPP and to \$1.90 using the new 2011 PPP.² Fixed in this way, the IPL simply corresponds to an arithmetic mean of the different national poverty lines for a given group of countries, all of which are weighted equally regardless of population size. It also assumes that countries in this group have common characteristics, meaning that differences among national poverty lines are random and cannot be explained by extra variables. An absolute poverty line corresponds to a given number of calories and to the cost of other objective necessary quantities, such as basic shelter, clothing, and health. If PPP is correctly established, the cost of the minimum basket of goods to satisfy the basic human needs in the least developed countries will be the same. In this paper, we shall call this group of countries the reference group, for which a common poverty line in PPP can be used.

However, the figures reported in the database of Ravallion *et al.* (2009) show that there is a relation between z_i and C_i for the reference group of very poor countries, even if not as close as for richer countries. For countries with a mean C_i lower than \$60 a month, the poverty line represents on average 92 percent of the mean consumption level, while it falls to 45 percent for the richer group of countries of the database. This last figure is much more in accordance with the usual definition of a relative poverty line, which is usually half the mean income in Europe (or 60 percent of the median income). We arrive at the first figure of 92 percent by computing the average of the reported poverty lines. For this average, we find \$38 a month with a standard deviation of \$12. How can such a large standard deviation be explained? In this group, the minimum and maximum poverty lines are \$19 and \$59, which means roughly between \$0.60 and \$2.00 a day. We can find two possible reasons, which are not exclusive. The first reason is the PPP calculations, as discussed in Deaton (2010). Our only answer is to assume that this could be partially compensated for by a less restrictive specification of the model, which would then relate the national poverty lines to a set of explanatory variables including mean consumption in the reference group. This is because national poverty lines reflect consumption patterns specific to each country, which may not be reflected with 2005 PPP or the 2011 PPP. In the empirical section, we shall propose the rate of unemployment as a possible indicator for official agencies when fixing their official poverty line. The second reason relates to the fact that, when fixing the national

²In fact, \$1.88, rounded up to \$1.90 for reporting convenience according to Ferreira *et al.* (2015).

poverty line, national agencies may be influenced by arguments which are partly subjective and related to social inclusion. In this case, a national poverty line does not consider simply the minimum level of subsistence, but may also partly be based on the amount of money needed to maintain a minimum acceptable way of living. We develop this aspect below, but we must first have a serious discussion about the data and the new 2011 PPP.

2.2. *The Database and its Limitations*

Our database comes partly from the appendix of Ravallion *et al.* (2009), who considered 74 developing countries. Their dataset includes national official poverty lines (PL) (or poverty lines computed by academics in some cases) and private consumption expenditure (PCE) per capita. The main sources are national accounts, as stated on page 169: “This article follows Ravallion, Datt, and van de Walle (1991) in using private consumption expenditure per capita from the national accounts.” These data refer to different years from 1988 to 2005 for different countries. They have been adjusted by the household consumption PPP collected during the International Comparison Program of 2005 (World Bank, 2008). The PCE and PL variables are reported on a monthly basis. This dataset is an improvement over the dataset previously used in Ravallion *et al.* (1991), which covered only 33 countries and had weaker price adjustment.³ We completed this initial database by adding several other variables taken from the World Bank website, including population size N_i , unemployment rate Ur_i , the Gini index, and poverty rates computed by the World Bank for two poverty lines (\$1.25, \$2.00). These extra variables were collected for the same dates as the initial data.⁴ If we decide to opt for the new 2011 PPP, as the World Bank did in computing its new IPL of \$1.90, we need four additional series. As explained in Ferreira *et al.* (2015), for a quantity expressed in 2005 PPP dollars and denoted by $x_{2005PPP}$, the conversion to 2011 PPP dollars, denoted by $x_{2011PPP}$, means that

$$(2) \quad x_{2011PPP} = x_{2005PPP} \times \frac{CPI_{2011}}{CPI_{2005}} \times \frac{PPP_{2005}}{PPP_{2011}},$$

where PPP_{2005} and PPP_{2011} are two PPP conversion series and CPI_{2011} and CPI_{2005} are the corresponding two series of consumer price indices. We have collected these series from the World Bank website.

The procedure used by the World Bank to compute the new IPL using 2011 PPP is well detailed and is discussed in Ferreira *et al.* (2015); namely, take the same database (the one we are using) and the same reference group of 15 countries as in Ravallion *et al.* (2009), convert the corresponding national poverty lines from 2005 PPP to 2011 PPP using equation (2), and take their mean over that group. Whereas Ravallion *et al.* (2009) obtained a \$1.25 IPL, this calculation

³Note, however, that this dataset is not exempt from oddities. For instance, the official poverty line for intermediate urban areas in Senegal was 661.7 CFA in 2005, \$1.06 at the current rate of exchange, while the use of PPP reduces the official poverty line to \$0.64.

⁴These new variables will help us for instance to estimate world income distribution around 2001, which is the average date of collection of the data in Ravallion *et al.* (2009).

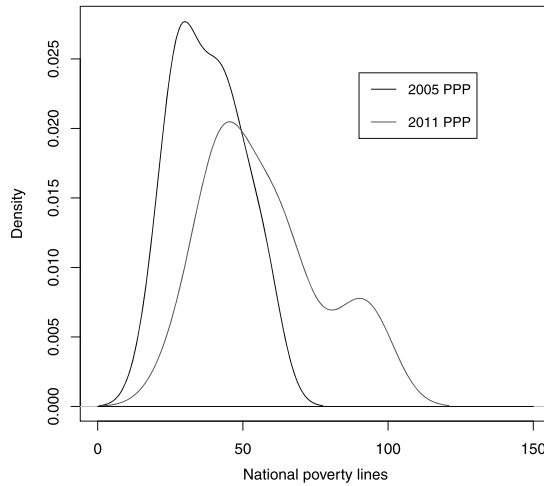


Figure 1. The impact of the change of PPP on the distribution of national poverty lines for the World Bank 15 countries’ reference group [Colour figure can be viewed at wileyonlinelibrary.com]

leads to \$1.88 when applying equation (2), a result that Ferreira *et al.* (2015) rounded up to \$1.90 for reporting convenience.

We see two immediate drawbacks to this updating approach. First, as demonstrated in Figure 1, the conversion formula introduces not only a drift in the distribution of the 15 converted national poverty lines of the World Bank reference group, a drift that would justify the rise to \$1.90, but also a major change in the shape of this distribution, with an increased variance and the appearance of a secondary mode. Second, the initial national poverty lines, reported in Ravallion *et al.* (2009), were in fact collected at various dates covering the range from 1985 to 2005 (80 percent of the dates are between 1992 and 2003). Due to this large range, it seems inadequate at the very least to apply the same conversion formula (equation (2)) to all these data.

These two remarks cast doubt on the neutrality of the conversion procedure. Doubts are also reported in Ferreira *et al.* (2015), who recommend use of the old \$1.25 and 2005 PPP for some countries for which the change implied by 2011 PPP would be too important. They eventually quote the option, once considered by the World Bank, which is simply “to not adopt the 2011 PPPs and continue to use the \$1.25 line in 2005 PPP-adjusted USD.”

As a final remark, Ferreira *et al.* (2015) explain that the \$1.25 IPL was coherent with the 2005 PPP, while the new \$1.90 IPL is now coherent with the 2011 PPP. They show that both poverty lines give similar rates of world poverty, because the change of PPP is meant to simply reflect a weaker value of the dollar when compared to most poorer countries’ currencies. They report a change in the global poverty rate from 14.5 percent to 14.1 percent for 2011 when passing from \$1.25 IPL and 2005 PPP to \$1.90 IPL and 2011 PPP. So even if this new IPL has created strong reactions in the press,⁵ it does not bring about a global change and

⁵According to the Financial Times (September 23, 2015): “The World Bank is to make the most dramatic change to its global poverty line for 25 years—raising its measure by a half to about \$1.90 per day—in a move likely to swell the statistical ranks of the world’s poor by tens of millions.”

it perhaps introduces more problems than it solves. For instance, it precludes comparisons with the existing literature, such as Sala-i-Martin (2006) or Deaton (2010). Therefore, we prefer to stay on the safe side and keep the data which are expressed in 2005 PPP, while waiting for a better database, which does not seem to be available for the time being. We shall nevertheless indicate in due course some variants of our results, variants obtained using the 2011 PPP.

2.3. Explaining Differences in National Poverty Lines

If the national poverty lines for the reference group do not simply reflect the cost of 2,100 calories per day per adult (as recommended by the World Bank; see Ravallion, 1994), this means that the national agencies are setting their poverty lines on partially subjective grounds. The poverty line thus gains positive elasticity with average consumption. There are various reasons for this. The composition of the reference basket of goods is socially determined. See the example given by Atkinson (1983, p. 188), where English workers went on strike because tea was planned to be withdrawn from the official basket of goods and replaced by milk, for computing the official poverty line. Despite the fact that tea has no nutritional value, it had a social value. The second reason is social inclusion, even in less developed countries. In the official Tendulkar, 2009 report on poverty evaluation in India, we find such a sentence as “Fundamentally, the concept of poverty is associated with socially perceived deprivation with respect to basic human needs.”

We now propose a small model which provides support for the inclusion of explanatory variables to determine the IPL. We borrow it from the literature on the subjective approach to poverty. We can draw a parallel between a national poverty line and the minimum income question (MIQ) that can be found, for instance, in Kapteyn *et al.* (1988). It is phrased as follows: “What is the minimum income that you would need in order to make the two ends meet?” If z_i is the reported answer, y_i the actual income of the household, and x_i a characteristic of the household (such as its composition), then the following regression⁶ is estimated over all the individuals i of a specific country:

$$(3) \quad z_i = \alpha + \gamma y_i + \beta x_i + \epsilon_i.$$

An estimated subjective poverty line for a given country corresponds to a fixed point (equating y and z) for every type of household composition x :

$$(4) \quad z^* = \frac{\hat{\alpha} + \hat{\beta}x}{1 - \hat{\gamma}}.$$

With a fixed point, we classify as poor those households who reported an income below the answer they gave to the MIQ, for each type of household composition.⁷

⁶A log-log regression is also possible, as in Van den Bosch *et al.* (1993).

⁷Other definitions and methods of establishing a subjective poverty line have been proposed in the literature; for instance, the Leyden poverty line as introduced in Goedhart *et al.* (1977), which is based on the Income Evaluation Question and the estimation of a particular social welfare function. Note also the CSP (Center for Social Policy) of the University of Antwerp introduced in Deleeck (1989) and reviewed in Van den Bosch *et al.* (1993), which relies both on the MIQ and on a supplementary question concerning financial ease. Only the households answering to that second question with some difficulty are taken into account in defining the poverty line.

Let us now return to the IPL and to our reference group. We assume that a national agency answers the minimum income question, revealing the value of the national poverty line. Now i denotes countries, and no longer individuals. Equation (3) becomes

$$(5) \quad z_i = \alpha + \gamma C_i + \beta x_i + \epsilon_i,$$

where z_i is the national official poverty line and C_i is the mean consumption level per capita for country i , x_i is a country-specific variable, and i covers the countries in the reference group. The same fixed-point algorithm provides the level of the poverty line, common to all countries of the reference group, but which can be a function of the country characteristics x_i :

$$z^* = \frac{\alpha + \beta x_i}{1 - \gamma}.$$

Ravallion *et al.* (2009) assumed that for the reference group, the restriction $\gamma = 0$ should be imposed, as well as $\beta = 0$. When these restrictions are imposed in equation (5), the common poverty line z^* becomes the empirical mean of z_i and its implicit elasticity with respect to national consumption is simply zero.

2.4. An International “Subjective” Poverty Line

As we first have to determine the composition of the reference group, an essential task is to test the validity of the model in equation (1) to see if we can really infer a consistent international poverty line from it. Throughout this paper, we use a Bayesian approach because equation (1) is a non-linear threshold regression model. In the classical approach, the estimator for the threshold parameter θ has a non-standard asymptotic distribution (for details, see Hansen, 2000), which makes it difficult to determine confidence intervals with the right coverage. The power of the Bayesian approach for this model was clearly shown in Bauwens *et al.* (1999, Ch. 8) and is detailed in Section 3 of this paper. One way of testing the validity of this model is to test for the effective presence of non-linearity. This can be done by testing the following assumption: $\alpha_1 = \alpha_2$ and $\gamma_1 = \gamma_2$. As shown in Bauwens *et al.* (1999, Ch. 8), the model is not identified under this assumption, so we have to introduce an informative prior to solve the problem. Estimating equation (1) under a non-informative prior for the regression coefficients, but a uniform informative prior for θ over the interval [32, 200] which amply covers the \$60 of Ravallion *et al.* (2009), we obtain posterior draws for all the parameters.⁸ Then we form the vector of differences δ :

⁸The precise Bayesian treatment for this model will be detailed in the next section. Here, we provide only the results of a test to justify our initial empirical model. The informative uniform prior for θ has a prior mean of 116 and a prior standard deviation of 48. With this prior, all values between 32 and 200 are equiprobable, but all values outside this range are excluded. We have chosen these values so as to cover a wide range while still identifying the model by keeping enough observations in each regime. We illustrate the shape of the posterior of θ in Section 4.

$$(6) \quad \delta = [\alpha_1 - \alpha_2, \gamma_1 - \gamma_2],$$

and draw the contours of these transformed draws. If zero belongs to a 90 percent confidence interval of δ , the model fails to be adequate. This is the case here. A Bayesian F -test (as reported in Bauwens *et al.*, 1999, Ch. 8, formula 8.55) has a marginal value of 0.27, which gives a p -value of 0.76 and clearly indicates that this model cannot be considered as non-linear. This model has already been criticized in the literature by Greb *et al.* (2011), among others. One of the results in the latter paper is that the model needs to be specified in logs; for example, by taking the consumption in logs. If we now consider this model, namely

$$(7) \quad z_i = s_i(\alpha_1 + \gamma_1 \log C_i) + (1 - s_i)(\alpha_2 + \gamma_2 \log C_i) + \epsilon_i,$$

and test linearity in the same way, this time linearity is rejected, since a Bayesian F -test of joint restriction has a posterior mean value of 277.7 and a p -value of 0.00. Consequently, we consider a model where consumption is expressed in logs instead of levels. Note that this kind of test would be difficult to implement in a classical approach such as Hansen (2000). In Hansen (2000), the conditional asymptotic distribution of the classical estimator of the vector of differences $\delta(\theta)$ defined in equation (6) is normal (conditional on a value of θ). But it would be very difficult to find the exact marginal asymptotic distribution of an estimator of δ . With our Bayesian approach, we have draws of the exact small-sample marginal distribution of δ .

The complete and extended model that we consider is as follows:

$$(8) \quad z_i = s_i(\alpha_1 + \gamma_1 \log C_i + \beta_1 x_i) + (1 - s_i)(\alpha_2 + \gamma_2 \log C_i) + \epsilon_i,$$

$$(9) \quad s_i = \begin{cases} 1 & \text{if } C_i < \theta, \\ 0 & \text{otherwise,} \end{cases}$$

$$(10) \quad \text{Var}(\epsilon_i) = s_i \sigma_1^2 + (1 - s_i) \sigma_2^2,$$

where z_i is an official poverty line in PPP dollars, $\log C_i$ the log of the average level of private consumption per capita in PPP dollars, x_i a set of explanatory variables specific to the first group, and θ is the unknown threshold. A different variance is allowed for each regime, because they correspond to two quite different mechanisms.

In this model, determination of the poverty line for the two groups is clearly based on different reasonings. A relative poverty line emerges for the richer group of countries, while the poverty line is based on a wide range of factors in the poorer group. Within this model, the new common poverty line for the reference group can be determined as a conditional expectation:

$$(11) \quad E(z_i | s_i = 1) = \alpha_1 + \gamma_1 E(\log C_i | s_i = 1) + \beta_1 E(x_i | s_i = 1).$$

In words, the poverty line that we propose for less developed countries is a function of a reference group consumption level, which is taken to be equal to an estimated fraction of the mean of the log consumption of that reference group and of

different contextual variables. It differs from the usual relative poverty line in that it depends not on national mean consumption per capita, but on the mean log consumption of a more general group, called the reference group. We call this new poverty line a “subjective” poverty line, not because it depends on subjective data, but for several other reasons. First, the implicit elasticity of this poverty line is neither zero nor one, as with the usual subjective poverty lines (see, e.g., Van den Bosch *et al.*, 1993). Second, as it depends both on the consumption level and the country characteristics x_i , it relates poverty to “inclusion in a particular society,” in the words of Atkinson and Bourguignon (2001). In the empirical section, we choose the unemployment rate as a measure of social inclusion. Third, our poverty line is defined with respect to a common group with which each country is supposed to identify. They may determine their poverty line by reference to that group. A final additional point in support of taking $\gamma_1 \neq 0$ concerns PPP. Taking the IPL as estimated only by α_1 is equivalent to assuming that the cost of consuming the necessary calories is the same across all the countries of the reference group. In other words, it amounts to saying that cost-of-living differences are perfectly equalized using either 2005 PPP or 2011 PPP. Figure 1 and the remarks made in Deaton (2010) show that this can be strongly questioned.

In Ravallion *et al.* (2009), and in most of the works coming from the World Bank as reported in Deaton (2010), the reference group is fixed. One criticism made by Deaton (2010) is that this creates discontinuity. For instance, revising PPP can remove a country (e.g. China) from the reference group and thus artificially increase or decrease the poverty line, thereby altering the number of poor people in the world. In our model, the reference group is determined endogenously and in a probabilistic way, which makes the problem of discontinuity less severe, especially since both regimes include the same variable $\log C_i$. We shall see, however, that the question of weighting remains.

2.5. *The Question of Weighting: China and India*

Discussing world poverty automatically poses the question of the weight of India and China. In our group of 74 countries, the total population is 4 billion 522 million (representing more than 70 percent of the world population at that time). India and China represent, respectively, 25 percent and 28 percent of this total. The remaining 47 percent are scattered among 72 countries. The weight of these two countries is considerable, as noted for instance by Deaton (2005), in computing poverty statistics. Deaton (2005) decided to present his results both for the weighted and the unweighted case. The argument goes as follows. On the one hand, if a common poverty line results from a political agreement at international institution level, then clearly each country should be given equal weight. If, on the other hand, the aim is to count poor people, clearly a weighted approach should be used. It should be noted, however, that weighted estimates also have some drawbacks, especially when weighting by population. This implicitly amplifies the institutional or technical bias when reporting national poverty lines for giant countries such as China and India. Weighting by population also gives too much weight to a large share of the population not suffering from poverty. For instance, following Deaton (2010), we could use only the number of poor people for weighting, instead

of the total population. This would give more weight to India, which has more poor people than China, although weighting by population gives more weight to China. Thus, we need to consider different weighting schemes.

In their approach, Ravallion *et al.* (2009) assume equal weighting for each country, because their revision of the IPL is simply an unweighted arithmetic mean of national poverty lines inside a given reference group. They choose as their reference group a group of countries having an average consumption per capita lower than \$60. The unweighted average IPL is \$1.25 and the weighted average is \$1.16 within that group. But with a threshold of \$60 a month, the reference group contains only 15 countries, and neither China nor India is a member. In the next section, we present Bayesian inference for our regression model with a threshold. We provide a posterior distribution for the threshold and compute the posterior probability of a country belonging to the reference group. This is a more realistic approach to the questions raised by weighting.

3. BAYESIAN INFERENCE FOR REGRESSION MODELS WITH A BREAK

The generic model we want to estimate is a two-regime regression model explaining z_i with a break determined when a variable C_i is lower or higher than an unknown threshold θ . It corresponds to one of the models described in Bauwens *et al.* (1999, Ch. 8), namely:

$$E(z_i|x_i) = x_i'\beta_1 \text{ if } C_i \leq \theta,$$

$$E(z_i|x_i) = x_i'\beta_2 \text{ if } C_i > \theta.$$

z_i is the dependent variable (national poverty lines), x_i a set of exogenous variables including a constant term, and C_i is the regime shift variable, which is supposed to be exogenous or predetermined. θ is a threshold parameter. We introduce the unobserved variable s_i , defined as follows:

$$s_i = \begin{cases} 1 & \text{if } C_i \leq \theta, \\ 0 & \text{otherwise.} \end{cases}$$

Regrouping these elements in a single equation, we obtain

$$z_i = s_i x_i' \beta_1 + (1 - s_i) x_i' \beta_2 + \epsilon_i,$$

where the error term ϵ_i is assumed to be normal with zero mean and constant variance σ^2 (the two-variance case will be treated below). For inference purposes, it is useful to define the following matrix:

$$(12) \quad X(\theta) = [s_i x_i', (1 - s_i) x_i'],$$

so that the model can be written in a more compact form:

$$(13) \quad z = X(\theta)\beta + \epsilon,$$

where z is a vector containing the N observations of z_i and β is the vector containing parameters β_1 and β_2 .

3.1. Likelihood and Posteriors

Considering N observations, the likelihood function of model (13) is as follows:

$$(14) \quad L(\beta, \sigma^2, \theta; z) \propto \sigma^{-N} \exp \left[-\frac{1}{2\sigma^2} \sum_{i=1}^N [z_i - X'_i(\theta)\beta]^2 \right].$$

Conditional on θ , this is the likelihood function of the usual regression model, so that natural conjugate prior densities for β and σ^2 belong to the normal inverted gamma2 family:

$$\begin{aligned} \pi(\beta|\sigma^2) &= f_N(\beta_0, \sigma^2 M_0^{-1}), \\ \pi(\sigma^2) &= f_{Ig}(\sigma^2 | v_0, s_0). \end{aligned}$$

A non-informative prior is obtained by letting the hyperparameters go to zero. The conditional posterior densities of β and σ^2 are as follows:

$$(15) \quad \pi(\beta|\theta, z) = f_t(\beta|\beta_*(\theta), s_*(\theta), M_*(\theta), v_*),$$

$$(16) \quad \pi(\sigma^2|\theta, z) = f_{Ig}(\sigma^2 | v_*, s_*(\theta)),$$

$f_t(\cdot)$ being the Student distribution. The different posterior hyperparameters are defined in the following way:

$$\begin{aligned} M_*(\theta) &= M_0 + X'(\theta)X(\theta), \\ \beta_*(\theta) &= M_*^{-1}(\theta)[X'(\theta)z + M_0\beta_0], \\ s_*(\theta) &= s_0 + \beta'_0 M_0 \beta_0 + z'z - \beta'_*(\theta)M_*(\theta)\beta_*(\theta), \\ v_* &= v_0 + N. \end{aligned}$$

The usual and convenient practice is to use a non-informative prior for the regression parameters with $\beta_0 = 0$, $M_0 = 0$, $s_0 = 0$, and $v_0 = 0$. But an informative prior on θ can be crucial. It is not possible to find a natural conjugate prior for the threshold parameter, so we are totally free to select this prior density as $\pi(\theta)$, with no further detail for the moment.

The posterior density of θ is obtained as a by-product of the Student posterior density in equation (15), being simply proportional to the inverse of its integration constant times the prior density of θ :

$$(17) \quad \pi(\theta|z) \propto |s_*(\theta)|^{-(N-k)/2} |M_*(\theta)|^{-1/2} \pi(\theta).$$

This posterior density does not correspond to any known form, and has to be analysed by numerical integration. In this case, a convenient choice for $\pi(\theta)$ is the

uniform distribution between bounds. The marginal posterior densities of β and σ^2 also have to be found using numerical integration, as we have

$$(18) \quad \pi(\beta|z) = \int f_t(\beta|\beta_*(\theta), s_*(\theta), M_*(\theta), v_*) \pi(\theta|z) d\theta,$$

and

$$(19) \quad \pi(\sigma^2|z) = \int f_{I_g}(\sigma^2|v_*, s_*(\theta)) \pi(\theta|z) d\theta.$$

The dimension of θ being one, we could use a traditional deterministic integration rule, like the Simpson rule, in order to evaluate these densities. However, as it is transformations of the parameters that interest us, a simulation method is better.⁹ As equation (17) is a marginal density, we just have to find a feasible grid over which to evaluate it, numerically compute the cumulative density, and then use the inverse transformation method to draw a value for θ , denoted by $\theta^{(j)}$ for $j = 1, \dots, M$. Briefly, the grid over which to evaluate equation (17) has to be chosen carefully, which means carefully selecting the bounds of the informative uniform prior. These bounds should cover most of the probability, but they should also avoid identification problems. As detailed in Bauwens *et al.* (1999, p. 235), the bounds should be chosen so as to ensure sufficient observations per regime. Then, we draw a value of θ from $\pi(\theta|z)$. Using this draw, we draw a value of β from the conditional posterior $\pi(\beta|\theta, z)$ given in equation (18), which is a Student density.

3.2. The Two-Variance Case

For modeling purposes, it is useful to consider the possibility of having different variances in the two regimes. If the endogenous variable is in levels and not in logs, then the variance of the error term is not scale-free. As the scale of the richer group of countries is larger, the variance of the error term should be larger. We cannot constrain the two regimes to have the same variance. We keep the same dichotomous variable s_i as in the original model and assume this time that

$$(20) \quad \text{Var}(\epsilon_i) = s_i\sigma_1^2 + (1 - s_i)\sigma_2^2.$$

Let us now set $\sigma_1^2 = \phi\sigma_2^2$, so that

$$\text{Var}(\epsilon_i) = \sigma_2^2(1 + s_i\phi - s_i) = \sigma^2 h_i(\theta, \phi),$$

as detailed in Bauwens *et al.* (1999, p. 236). Let us now scale the observations by $\sqrt{h_i(\theta, \phi)}$ in order to obtain a regression model with homoskedastic errors of variance σ^2 :

⁹It is very easy to compute the posterior density of a transformation of a parameter when we have posterior draws from this parameter. We just have to take the transformation of each draw as a draw from the posterior of the transformed parameter. The use of deterministic integration rules leads to much more complicated procedures.

$$(21) \quad z(\theta, \phi) = [z_i / \sqrt{h(\theta, \phi)}],$$

$$(22) \quad X(\theta, \phi) = [s_i x'_i / \sqrt{h(\theta, \phi)}, (1 - s_i) x'_{ii} / \sqrt{h(\theta, \phi)}].$$

The regression model becomes

$$z(\theta, \phi) = X(\theta, \phi)\beta + \epsilon,$$

its likelihood function being

$$(23) \quad L(\beta, \sigma^2, \theta, \phi; z) \propto \sigma^{-N} \prod_{i=1}^N h_i(\theta, \phi)^{-1/2} \\ \times \exp \left[-\frac{1}{2\sigma^2} (z(\theta, \phi) - X(\theta, \phi)\beta)' (z(\theta, \phi) - X(\theta, \phi)\beta) \right].$$

The prior densities on β , σ^2 , and θ are the same as before. We have to introduce a new prior for ϕ , namely $\pi(\phi)$. The conditional posterior densities of β and σ^2 have the same form as before. We just have to replace z by $z(\theta, \phi)$ and $X(\theta)$ by $X(\theta, \phi)$ in the necessary expressions. The joint posterior density of θ and ϕ is as follows:

$$(24) \quad \pi(\theta, \phi|z) \propto \prod_{i=1}^N h_i(\theta, \phi)^{-1/2} |s_*(\theta, \phi)|^{-(N-k)/2} |M_*(\theta, \phi)|^{-1/2} \pi(\theta)\pi(\phi).$$

It is slightly more difficult to draw $\theta^{(j)}$ and $\phi^{(j)}$ jointly from this bivariate density equation (24) than it is to draw $\theta^{(j)}$ from the univariate density equation (17). A feasible method can be found if we remember that it is always possible to decompose a bivariate density into the following:

$$\pi(\theta, \phi|z) = \pi(\phi|\theta, z) \times \pi(\theta|z).$$

Consequently, we first draw from the marginal density $\pi(\theta|z)$ and then from the conditional $\pi(\phi|\theta^{(j)}, z)$. To apply this method, we first need to determine a grid over θ and ϕ in order to fill up a matrix. From this matrix of points, we can numerically determine the marginal density $\pi(\theta|z)$. For a given draw $\theta^{(j)}$, we have to find the corresponding conditional $\pi(\phi|\theta^{(j)}, z)$. Of course, we will not have a draw $\theta^{(j)}$ that corresponds exactly to a line of the initial matrix of points. So we shall have to proceed by linear interpolation between two lines, as explained in the Appendix (in the Online Supporting Information).

3.3. Weighting by Population

Let us now consider the case where we want to weight the regressions using the population or the number of poor people as weight w_i . We simply have to change the definition of $h_i(\theta, \phi) = 1 + s_i(\phi - 1)$ in equation (20) into the following:

$$h_i(\theta, \phi) = \frac{w_i}{sw_2} + s_i(\phi * \frac{w_i}{sw_1} - \frac{w_i}{sw_2}),$$

where sw_1 and sw_2 represent the average of w_i in regimes one and two, respectively, so that we recover the original formula in the unweighted case.

3.4. Posterior Distribution of the IPL

Now, we need to find the posterior density of the *IPL*, based on the first-regime characteristics. It is obtained as a transformation of the parameters and of an evaluation of the average characteristics of the reference group taken conditionally on the draws of θ . In the empirical section, we test and accept the restriction $\alpha_1 = 0$ in the first regime when there are extra exogenous variables. The new *IPL* with this more parsimonious model is as follows:

$$(25) \quad E(z_i | s_i = 1) = \gamma_1 E(\log C_i | C_i < \theta) + \beta_1 E(x_i | C_i < \theta).$$

If we now take into account weights w_i , expectation $E(z_i | s_i = 1)$ becomes

$$(26) \quad E(z_i | s_i = 1) = \gamma_1 E(w_i \log C_i | C_i < \theta) + \beta_1 E(w_i x_i | C_i < \theta).$$

These two quantities are functions of the posterior density of γ_1 , β_1 and θ .

We can obtain draws of the posterior density of the *IPL* in the following way. We first draw $\theta^{(j)}$ and $\phi^{(j)}$ from the joint posterior density equation (24). We then determine a sample separation. Conditional on this sample separation, we compute a possibly weighted sample mean for variables $\log C$ and x . We then draw $\beta_1^{(j)}$, $\gamma_1^{(j)}$, and $\sigma^{(j)}$ from their conditional posterior densities $p(\beta_1, \gamma_1 | \theta, \phi, \sigma^2, z)$, which is a conditional normal density, and $p(\sigma^2 | \theta, \phi, z)$, which is a conditional inverted gamma2. By combining these draws and sample means, we obtain a draw from the posterior density of the *IPL*. Once we have enough draws, we can compute a mean and a standard deviation, and plot a posterior density. Formally, a draw $z^{(j)}$ corresponds to a transformation of $\gamma_1^{(j)}$, $\beta_1^{(j)}$, $\theta^{(j)}$, and $\sigma^{(j)}$:

$$(27) \quad z^{(j)} = \gamma_1^{(j)} \sum_{i=1}^n w_i \log(C_i) 1(C_i < \theta^{(j)}) + \beta_1^{(j)} \sum_{i=1}^n w_i x_i 1(C_i < \theta^{(j)}),$$

where the w_i are weights (equal or unequal) summing to one according to the scheme $\sum_i w_i \mathbb{1}(C_i < \theta^{(j)}) = 1$. The unweighted case corresponds to $w_i = 1/n_{1i}$, where n_{1i} is the number of observations in the first regime given the j th draw.

4. EMPIRICAL COMPARISON OF DIFFERENT BAYESIAN POVERTY LINES

We now present our empirical results for three different cases. First, we present the unweighted case, which amounts to using raw data. Then, we use weights: either the population as suggested in Deaton (2005), or the number of poor people below the official poverty line as suggested in Deaton (2010).

4.1. *The Two-Regime Model*

The most general model that we start with is as follows:

$$(28) \quad z_i = s_i(\alpha_1 + \gamma_1 \log C_i + \beta_1 Ur_i) + (1 - s_i)(\alpha_2 + \gamma_2 \log C_i + \beta_2 Ur_i) + \epsilon_i,$$

using two variances for the error term, depending on the regime. This model fits the notion that variable Ur_i , the unemployment rate, can help to predict the poverty line in the reference group, under the intuition that a higher rate of unemployment would lead to a higher official poverty line.¹⁰ Investigating the unemployment variable's significance in the model might shed light on the varying national poverty lines found in the reference group of Ravallion *et al.* (2009).

Using a uniform prior for θ over the range [80,200], a uniform prior on ϕ over the range [0.001, 0.25] and non-informative priors over the other parameters, we conduct a specification search, first with no population weighting and then weighting either by population or by number of poor people. We consistently reach the same specification, displayed in Table 1.

The first regime requires the presence of both $\log C_i$ and an extra variable to explain the level of the national poverty lines, while the constant term plays no role. In contrast, the second regime has the form of an affine function in $\log C_i$, with no other explanatory variables. The model is clearly non-linear, first because the value $\phi = 1$ does not belong to a 90 percent posterior confidence interval of ϕ ; respectively, [0.042, 0.125], [0.034, 0.114], and [0.012, 0.038] for the three approaches to weighting. So the two error term variances can never be the same in the two regimes. Second, the two γ 's are statistically different, as no credible posterior confidence interval of their difference could contain the value 0.0. As a matter of fact, 90 percent posterior confidence intervals for $\gamma_1 - \gamma_2$ are equal, respectively, to [-103.7, -97.6], [-105.8, -98.3], and [-117.6, -104.2] with the three approaches to weighting. Thus, two regimes really are needed.

The posterior density of θ is displayed in the three panels of Figure 2. It is unimodal in the unweighted case, but presents slight secondary modes in the two weighted cases. Weighting has a strong influence on the position of the modes. The average consumption level needed to determine the upper bound of the reference group (in fact, the posterior expectation of θ) is \$169 in the unweighted case. This is nearly three times the level of \$60 chosen in Ravallion *et al.* (2009). This latter value does not belong to a 90 percent posterior confidence interval of θ , which is [\$139, \$182]. For the weighted cases, the posterior expectation goes down to \$140 when weighting by population, but goes up to \$174 when weighting by the number of poor people. These values are still a long way from \$60, which is still not contained in a 90 percent confidence interval (respectively, [\$125, \$180] and [\$135, \$182]). The sample separation is not much influenced by these changes, as the positions of China and India are not

¹⁰Unemployment as a percentage of the total labor force, probably only covering the formal sector. The source is the World Bank website.

TABLE 1
EXPLAINING NATIONAL POVERTY LINES USING A TWO-REGIME MODEL WITH TWO VARIANCES,
2005 PPP

	Unweighted		Population Weighted		NBER Poor Weighted	
	Mean	SD	Mean	SD	Mean	SD
First regime						
γ_1	8.64	(0.72)	10.26	(0.59)	10.30	(0.61)
U_r	0.92	(0.16)	0.78	(0.06)	0.72	(0.07)
Second regime						
Intercept	-496.0	(146.6)	-505.6	(185.9)	-520.1	(323.9)
γ_2	109.3	(1.87)	112.5	(2.25)	119.8	(3.83)
θ	169.2	(14.03)	140.0	(17.04)	174.2	(12.88)
σ_1^2	197.8	(47.36)	1739.0	(505.4)	1386.8	(346.7)
σ_2^2	2767.4	(741.4)	27989.6	(7149.4)	64792.1	(16943.2)
ϕ	0.076	(0.026)	0.066	(0.025)	0.023	(0.008)
w_{China}	1/74=0.014		0.28		0.23	
w_{India}	1/74=0.014		0.25		0.26	

Figures correspond to posterior mean and posterior standard deviation. The last two lines indicate the average weights given to China and India in the three different weighting schemes.

greatly affected. The biggest change occurs with Indonesia. Any attempts to increase the prior range of θ lead to the same results.

The posterior probability of belonging to the reference group is evaluated during the Monte Carlo sampling by counting the number of times the condition $C_i \leq \theta^{(j)}$ is verified. Table 2 gives information on that probability, together with a list of the major countries. Weighting or not weighting affects the composition of this group, but does not affect the position of the two major countries, China and India. Weighting by population moves Indonesia out of the reference group by lowering its probability of belonging to 0.17.

The two very large countries, China and India, have very low national poverty lines (\$26 for China and \$27 for India, per month). Consequently, when more weight is put on these countries, the value of γ_1 increases from 8.64 to 10.26 or 10.31. Weighting therefore has a strong effect on the posterior density of γ_1 . This will greatly affect our modeled international poverty line.

4.2. Modeling the Poverty Line

What is the real influence of population weighting when determining the IPL? If we simply compute the mean of the national poverty lines inside our reference group, we obtain the result as a by-product of the Monte Carlo integration. Using this sample-determined reference group, we compute two different means, an unweighted mean and a population weighted mean. In the first case, we obtain \$1.48 (0.020), while in the other case we obtain \$1.01 (0.007). So weighting by population (or by the number of poor) leads to a lower poverty line when we compute it as a mean, whatever the method of weighting. Weighting had the same effect when considering the small 15 country reference group of Ravallion *et al.* (2009).

Let us now report our inference results concerning our modeled IPL.

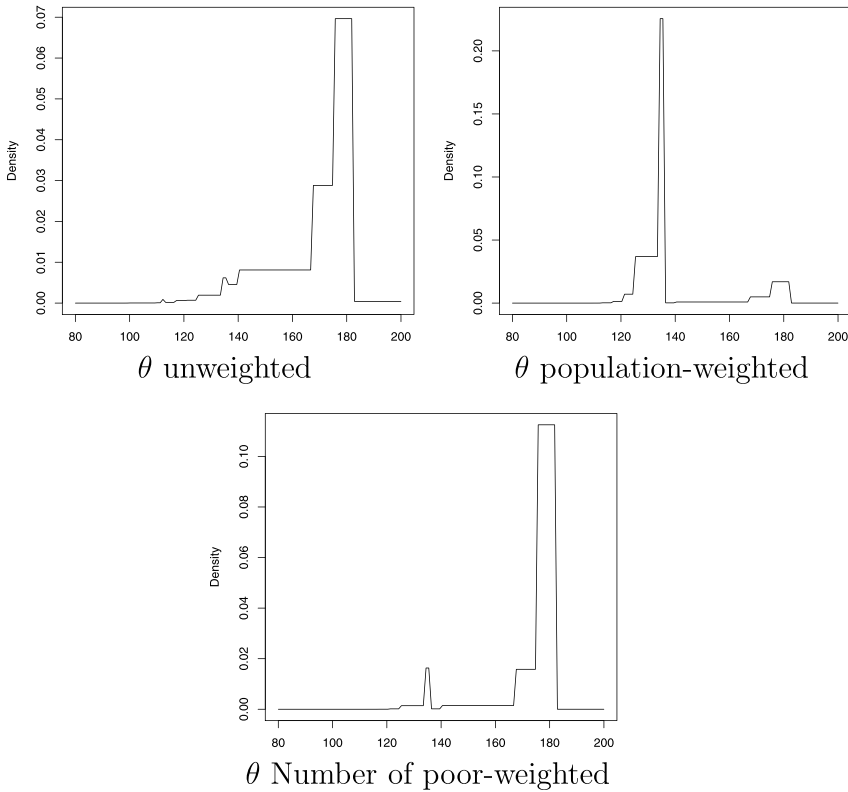


Figure 2. The posterior density of θ for the unweighted and weighted cases

Equations (25) and (26) define the poverty line as a linear function of the average log consumption and the average unemployment rate taken inside the reference group and possibly weighted by population share. This time, the weighting scheme influences both the values of the regression coefficients and the mean value of the regressors. With the first model, Table 3 reports an IPL of \$1.48 (0.036), which leads to a 90 percent posterior confidence interval of [\$1.30, \$1.65]. With the second model, the population-weighted regression, the IPL increases to \$1.65 (0.085) with a 90 percent posterior confidence interval of [\$1.50, \$1.79]. Weighting by the official number of poor people yields nearly the same IPL, \$1.63 (0.088) and a 90 percent posterior confidence interval of [\$1.49, \$1.78].

Remark. *As a side remark, if we had used the 2011 PPP, we would have obtained a modeled poverty line of \$2.29 in the unweighted case, of \$2.39 in the population weighted case, and of \$2.46 if weighting by the official number of poor, and hence again a greater value than the \$1.90 IPL of the World Bank. But these figures have to be taken with a grain of salt in accordance with the remarks we made concerning the 2011 PPP conversion.*

TABLE 2
THE PROBABILITY OF BELONGING TO THE REFERENCE GROUP

Case	$Pr = 1.0$	$0.9 < Pr < 1.0$	$Pr = 0.0$
Unweighted	30	9	32
	Bangladesh	China	Brazil
	India	Indonesia	Mexico
	Pakistan		Russia
Population weighted	34	2	33
	Bangladesh	China	Brazil
	India		Mexico
	Pakistan		Russia
NBER poor weighted	36	3	32
	Bangladesh	Indonesia	Brazil
	China		Mexico
	India		Russia
	Pakistan		

Countries with more than 100 million inhabitants are listed. The numbers indicate the average size of the group.

So, whatever the method, without population weighting we have a fairly unvarying revised IPL of around \$1.48, using 2005 PPP. However, if weighting changes things, its effect depends on the method we use to compute the IPL. Weighting lowers the IPL when it is calculated from a reference mean. This is because large countries such as Bangladesh, China, and India have official poverty lines which are far below \$1.25. Weighting has an inverse effect when the IPL is drawn from the parameters of our regression model, leading to an increase in its value. This is not due to the difference between the raw and weighted means of our regressors but, rather, to the increase in γ_1 . Deaton (2005) says that there are as many good reasons for weighting as there are for not weighting, and concludes that we should just provide both results. We shall follow his advice.

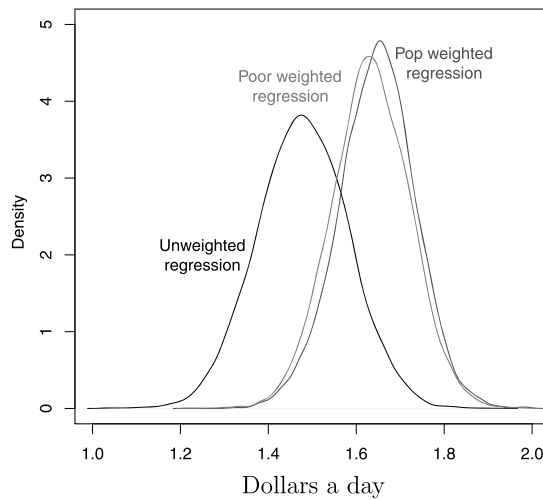


Figure 3. The posterior density of the modeled IPL using 2005 PPP [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 3
WHICH INTERNATIONAL POVERTY LINE?

Regression Model	IPL	IPL
	2005 PPP	2011 PPP
Unweighted regression	1.48 (0.036)	2.29 (0.218)
Population weighted regression	1.65 (0.085)	2.39 (0.166)
NBER poor weighted regression	1.63 (0.088)	2.46 (0.161)

The second column provides the posterior mean and standard deviation of our modeled IPL for each model in column one, using 2005 PPP. Draws for the modeled IPL are obtained using equation (27). The last column in italics corresponds to 2011 PPP and is reported only as an indication.

We now measure the influence of these different IPLs on the attempts to determine the number of poor people in the world. This requires us to propose a way to estimate the world income distribution. We also discuss another way of determining the IPL as proposed in Deaton (2010, p. 17).

5. HOW BEST TO ESTIMATE THE WORLD INCOME DISTRIBUTION?

To get some idea of how the way we determine the poverty line affects the number of poor people in the world and their location, we need a tool to estimate the world income distribution—or at least the overall income distribution of our group of countries, for which the IPL is designed. The introduction discusses the various options described in the literature, and we now present our contribution.

5.1. A Parametric Income Distribution for the Developing World

We build on the idea—found, for instance, in Atkinson and Bourguignon (2001)—that in each country, the income distribution can be represented by a lognormal density $f_{\Lambda}(y|\mu, \sigma^2)$, enabling two parameters to be estimated for each country. We know that the mean of a lognormal distribution is equal to $\exp(\mu + \sigma^2/2)$. Atkinson and Bourguignon (2001) propose to calibrate the mean using the PPP figure for daily consumption per capita, while pegging σ^2 at different prior values. Holzmann *et al.* (2007) propose to calibrate σ^2 using data on the Gini coefficient, as the formula for the Gini in the lognormal density is $2\Phi(\sigma/\sqrt{2}) - 1$, which depends only on σ . Data for the Gini coefficients are available from the World Bank, presumably using different sources, such as survey data. In both papers, the world income distribution is obtained by aggregating national income distributions using population shares.

These two types of information (consumption per capita and Gini coefficient) may not be sufficient to provide a precise indication of the shape of the left tail of the distribution if uncertainty concerning the value of the Gini coefficient is high. For instance, the Gini for consumption and the Gini for income might not be the same. The World Bank provides extra information that can be used to model the left tail of the income distribution, in the form of headcount poverty rates for two values of the poverty line, namely \$1.25 and \$2.00, using 2005 PPP.

The theoretical headcount poverty rate corresponds to $F_{\Lambda}^{-1}(1.25|\mu, \sigma)$ for a \$1.25 poverty line, for instance, where $F_{\Lambda}^{-1}(\cdot)$ represents the quantile function of the log-normal distribution. We have collected these supplementary data in order to construct a loss function for each country:

$$\begin{aligned} \text{loss} = & (pv_{1.25} - F_{\Lambda}^{-1}(1.25, \mu, \sigma))^2 + (pv_{2.00} - F_{\Lambda}^{-1}(2.00, \mu, \sigma))^2 \\ & + (C - \exp(\mu + \sigma^2/2))^2 + (Gi - 2 * \Phi(\sigma/\sqrt{2}) + 1)^2. \end{aligned}$$

Here, $pv_{1.25}$ is the empirical headcount for \$1.25 and $pv_{2.00}$ the corresponding value for \$2.00. C is the empirical daily mean consumption per capita, and Gi is the empirical Gini coefficient for one country.¹¹ This is a method of moments. We propose to minimize this loss function for each country separately. We could, however, have tried joint minimization, introducing common measurement errors to take into account the different origins of the data. This is certainly a limitation of our approach.

Nevertheless, we have managed to minimize our loss function for each of the 74 countries of our sample with no significant outlier. Using this method, we aim to obtain an income distribution for each country that is consistent with both the macro data of mean consumption per capita and with some microeconomic measures of dispersion, in particular for the left tail of the distribution. Then we aggregate these national adjusted distributions using population pop_i as a weight, imposing that the weights sum to one to obtain the world distribution of income (WDI) $f_W(y)$:

$$f_W(y) = \sum_{i=1}^{74} w_i f_{\Lambda}(y_i|\mu_i, \sigma_i^2), \quad w_i = pop_i / \sum pop_i.$$

Figure 4 shows the graph of this estimated mixture of 74 lognormals, together with two poverty lines, the old \$1.00 a day and our revised proposal of \$1.48 (without weighting).

China and India represent 53 percent of the population of our sample. They have different income distributions: China is richer, but with more inequality. The overall distribution is fairly smooth, probably because we have only 74 countries representing the developing world. Very rich countries or regions such as the United States and Europe are excluded, so no income polarization at world level can be detected.

5.2. Deaton's Alternative Poverty Line

Using our estimation of the WDI, we can now try to implement the alternative method suggested in Deaton (2010, p. 17) for computing a common poverty line as a weighted mean of all national lines. The argument put forward in Deaton (2010) is that we need to get rid of the discontinuity created by the point estimate of the reference group. Discontinuity refers to the fact that the exit of China or India from the reference group would induce a sharp rise in the mean reference poverty level. The

¹¹An update using the 2011 PPP would, of course, mean considering a totally new dataset, with updated Gini coefficients and new poverty headcounts for \$1.90 and \$3.10, which are the two values of the poverty now documented in Povcal.

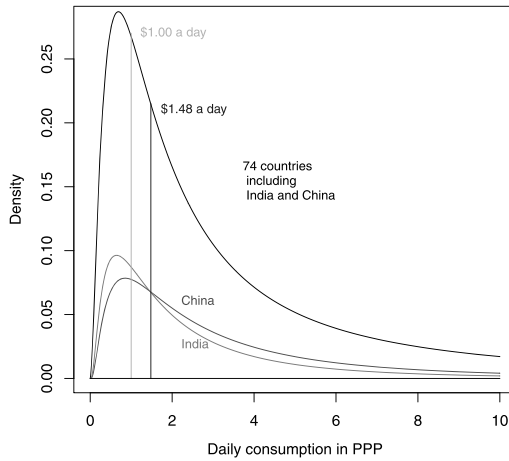


Figure 4. The income distribution for 74 developing countries around 2001 [Colour figure can be viewed at wileyonlinelibrary.com]

weights have to be determined by the implied poverty rates of each country, which themselves are determined by the common poverty line inserted in the expression of the cumulative distribution $F_{\Lambda}(\cdot)$ of the income distribution inside each country. So it is a fixed-point method that needs to be implemented. This becomes easy once we have an estimate of the income distribution for each country. Formally, we have

$$(29) \quad \bar{z} = \sum_{i=1}^n z_i w_i(\bar{z}), \quad w_i(\bar{z}) = \frac{F_{\Lambda}(\bar{z} | \hat{\mu}_i, \hat{\sigma}_i)}{\sum_{i=1}^n F_{\Lambda}(\bar{z} | \hat{\mu}_i, \hat{\sigma}_i)}.$$

\bar{z} has to be found iteratively, once all the μ_i, σ_i have been estimated. There is still, of course, the question of defining the weights. In formula (29), population is not taken into account in the definition of w_i . Deaton (2010) recommends instead using a formula where the weights include population:

$$(30) \quad w_i = \frac{\text{Pop}_i F_{\Lambda}(\bar{z} | \hat{\mu}_i, \hat{\sigma}_i)}{\sum_{i=1}^n \text{Pop}_i F_{\Lambda}(\bar{z} | \hat{\mu}_i, \hat{\sigma}_i)}.$$

In the first case, the weights are a function of the headcount ratio. In the second case, the weights depend on the number of poor people and thus on population size. This makes a huge difference, because in the second case the main weight is given to India and China, just because of their huge populations.

Using iterations and equation (29), the common poverty line is found to be \$1.64. However, when population is included in the definition of weights, as in equation (30), this goes down to \$1.15. In fact, the method suggested by Deaton

(2010) reacts in the same way as when population weighting is introduced into the mean method based on a reference group.

If we adopt a common poverty line for the 74 countries of the database, the poverty line suggested by Deaton with equation (30) leads to a figure of 1,172 million poor people, while if we use equation (29), this figure goes up to 1,672 million.¹² In the next section, we follow Atkinson and Bourguignon (2001) by arguing that two poverty lines are needed to measure the number of poor people in the world—one for the reference group and another for the other group—because they are not based on the same rationale.

6. THE POVERTY COUNT IN THE DEVELOPING WORLD

We are now equipped with all we need to obtain a posterior distribution of the number of poor people. We have a probabilistic way of determining the composition of the reference group, the distribution of the IPL, and a representation (albeit imperfect) of the income distribution of each of the countries. First, we characterize the number of poor people inside the reference group and then we generalize these computations to the whole group of 74 countries in our database.

6.1. Modeling the Poverty Count in the Reference Group

We have provided a proper way of defining the group of countries for which a single poverty line can be used, the IPL for which we have a posterior density. What is the number of poor people in that group? We shall work conditionally on our estimated WDI. First, for each draw of the parameters of our two-regime model, we obtain a sample separation depending on the value of $\theta^{(j)}$. The countries in the reference group are those for which $C_i < \theta^{(j)}$. We then deduce a draw for the poverty line $z^{(j)}$. For each country of the reference group, we compute a poverty rate by inverting its lognormal income distribution:

$$h_i^{(j)} = F^{-1}(z^{(j)} | \hat{\mu}_i, \hat{\sigma}_i^2).$$

We then multiply this rate by the national population N_i to obtain the corresponding number of poor people in that country. By aggregation over the countries of the reference group, we obtain a draw for the posterior density of the number of poor people in the reference group:

$$np^{(j)} = \sum_{i \in [C_i < \theta^{(j)}]} h_i^{(j)} N_i.$$

For M draws, we have an estimation of the posterior density of the number of poor people in the reference group which takes into account the stochastic

¹²For each country, we can obtain the poverty headcount h_i for a given value of the poverty line just by inverting its lognormal distribution. We then compute the number of poor as the product $\text{poor}_i = h_i \times N_i$, where N_i is the total population in that country. We obtain the total number of poor people by summation over the 74 countries.

composition of that group. In the left panel of Figure 5, we provide a graph of our posterior density for three different models of the IPL.

With the unweighted model, we obtain a number of poor people averaging to 1,448 million (and a standard deviation, SD of 105). If we weight by population, the average number of poor people rises slightly to 1,505 million (SD of 89). Weighting the regression by the official number of poor people increases the mean slightly to 1,584 million (SD of 78). These figures, presented in Table 4, differ greatly according to the method of weighting. This is mainly because we are only focusing on the reference group, whose composition changes with the method of weighting. There is a strong discontinuity effect. The rightmost curve corresponds to the highest number of poor people derived from a mean poverty line of \$1.63, while the curve in the middle is derived from a slightly higher poverty line (\$1.65) but indicates a slightly lower number of poor people. This is simply because the type of weighting chosen affects whether or not China and Indonesia are included in the reference group, as seen from Table 2. We are thus looking for a poverty line definition and an evaluation of the number of poor people that would be less sensitive to discontinuity.

6.2. Modeling the Poverty Count in the Developing World (74 Countries)

The objective of Deaton (2010) was to find a mechanism to determine an international poverty line which did not include any discontinuity. That meant not having a reference group. However, the argument led by Atkinson and Bourguignon (2001) supports having at least two different poverty lines, depending on the income level of the different countries: "... to provide a framework which unifies the measurement of poverty in developing and developed countries." In this subsection, we try to simultaneously analyze the determination of a poverty line and the determination of the number of poor people for our 74 countries, representing most of the developing countries and a sample of moderately developed countries.

Our two-regime model assumes two types of poverty lines, an IPL common to all the countries in the reference group and a collection of relative poverty lines, each specific to a country outside this group. For this second poverty line, we can take the national poverty line, provided that it is greater than the random draw of

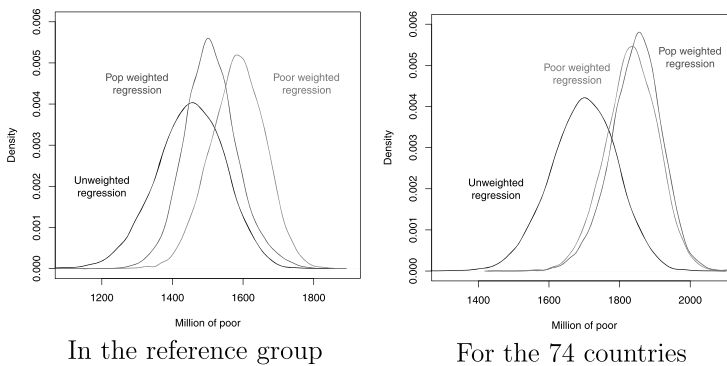


Figure 5. The posterior density of the number of poor people around 2001 [Colour figure can be viewed at wileyonlinelibrary.com]

TABLE 4
THE POVERTY COUNT IN THE DEVELOPING WORLD AROUND 2001 (MILLIONS)

Group Poverty Line	Reference IPL	World $\max(IPL, z_i)$	China IPL	India IPL
Unweighted	1,448 (105)	1,698 (95)	409 (32)	498 (32)
Population weighted	1,505 (89)	1,846 (72)	459 (25)	547 (24)
Poor weighted	1,584 (78)	1,833 (75)	455 (26)	543 (25)
Official figures	1,195	1,599	360	416

Official figures were computed using the official poverty rate at the national poverty line. No figures exist for 43 countries in the World Bank dataset, so we determined which of the normalized poverty lines of the World Bank (\$1.25, \$2.00, \$2.50, \$4.00, and \$5.00) was closest to the national poverty line and took the corresponding poverty rates.

the IPL. So for each draw of the parameters of our two-regime model, we have the following:

$$(31) \quad \tilde{z}_i^{(j)} = \mathbb{1}(C_i \leq \theta^{(j)}) IPL^{(j)} + \mathbb{1}(C_i > \theta^{(j)}) \max(z_i, IPL^{(j)}).$$

From this draw $\tilde{z}_i^{(j)}$, we determine 74 poverty rates:

$$h_i^{(j)} = F^{-1}(\tilde{z}_i^{(j)} | \hat{\mu}_i, \hat{\sigma}_i^2),$$

which are aggregated into

$$np^{(j)} = \sum_{i=1}^{74} h_i^{(j)} N_i,$$

in order to obtain a draw from the posterior density of the number of poor people in the world. This procedure involves no specific discontinuity but, rather, a comprehensive definition of the number of poor people, mixing both absolute poverty and inclusion. The right panel of Figure 5 contains three graphs of this posterior density, depending on the method of weighting. There is still a difference between weighting or not weighting, but the method of weighting is now less of a factor. Moreover, the ordering of the posterior densities of the number

TABLE 5
THE LOCATION OF POOR PEOPLE IN THE DEVELOPING WORLD AROUND 2001

Region	Unweighted	Population Weighted	Poor Weighted	Sala-i-Martin
Africa	245	263	262	297
East Asia	576	639	634	41
East. Europe	36	36	36	4
Latin America	177	177	177	21
MENA	26	29	28	1
South Asia	639	702	697	33

Figures are in millions. It was not possible to obtain feasible standard deviations because the poverty line is fixed outside the reference group. The last column comes from Table 2 of Sala-i-Martin (2006), based on a \$1.50 a day poverty line.

of poor people is now consistent with the ordering of the level of the mean poverty lines.

Using the comprehensive definition of a poverty line given in equation (31), we reach an estimation of the number of poor people in the developing world, reported in Table 4, corresponding to a period around 2001. Table 5 details how these individuals are divided among the six traditional regions of the World Bank.

Most of the poor are located in East Asia (China) and South Asia (India). Weighting does not have much of an influence on the ranking of poverty counts. These figures are much higher than in Sala-i-Martin (2006), who finds around 400 million poor people in 2000 using a common poverty line of \$1.50. Using our IPL defined in equation (31) (which is \$1.48 in the unweighted case for the reference group), we find 1,698 million. Our figures for Africa are similar to those of Sala-i-Martin (2006). Where we find hugely different figures is for East and South Asia, essentially China, India, and Indonesia. It is worth comparing the data themselves. We collected the official poverty rates from the website of the World Bank. When we multiply the official poverty rates by the population and sum up the countries, we find a total of 1,599 million poor people. Our unweighted evaluation of 1,698 million appears to be consistent with the information contained in our database. Does this raise any questions about the way in which Sala-i-Martin (2006) estimated his WDI? Or is it simply a problem of data reliability? We provided posterior densities and posterior confidence intervals. These are, of course, conditional on the available data and we did not include in our estimation the possibility of measurement error. However, beyond the question of measurement error, different types of data lead to different evaluations of consumption and inequality. More precisely, there can be huge differences between consumption evaluated using national account data and consumption evaluated using survey data (see, e.g., Deaton, 2005). Our WDI is calibrated so as to reproduce national account-based measurement of consumption per capita, whereas Sala-i-Martin (2006) makes extensive use of survey data.

7. CONCLUSION

This paper yields three findings. The first concerns the revision of the IPL. Using the same data as Ravallion *et al.* (2009), we provide a consistent model leading to a revision of the IPL from \$1.25 to a value between \$1.48 to \$1.65, depending on the weighting scheme, and providing confidence intervals and posterior densities. If we had used the new 2011 PPP conversion, these figures would have been largely increased, and in any case greater than \$1.90. Whatever the PPP, the World Bank always underestimates the IPL because it neglects social inclusion. The second finding is that weighting strongly affects the final result, and that the change depends greatly on the model specification for deriving the IPL. The third finding concerns the number of poor people in the world—at least, the world as represented by the 74 countries of the database around 2001. Our evaluation varies between 1,698 million and 1,846 million. It relies on a fairly simple evaluation of the WDI which is realistic and in agreement with the official

figures collected from the World Bank. Of course, all our results are conditional on the data that we used. Possible limitations include errors of observation, the fact that observations correspond to different periods, and the fact that the WDI was estimated from only a few moments and quantiles. We did not introduce any provision in our model for these factors. Thus, the usual caution should be applied when interpreting these results.

There are three main messages in our findings. First, an IPL is not simply the price of 2,100 calories per day adjusted by PPP. It has to take into account local characteristics and is affected to some extent by average consumption in the country and by social inclusion. The second message is that when counting poor people in the world (or in a significant part of the world), the IPL cannot be used outside what we call here a reference group, because outside the reference group, poverty is no longer a matter of survival but also, and mainly, of social inclusion. Moreover, our model clearly shows that national poverty lines outside the reference group are a function of the country's mean consumption. The final message concerns the data, their quality, and their sources. We calibrated our WDI on national accounts (coming from Ravallion *et al.*, 2009) and found that there were 576 million poor people in East Asia based on our \$1.48 poverty line. Using PovcalNet from the World Bank with \$1.50 a day and 2005 PPP, we found for 2002 a headcount of 36.5 percent and 692 million poor people. Using survey data, 2005 PPP, and \$1.50, Sala-i-Martin (2006), we found only 41 million for the whole of East Asia. This clearly raises questions about the implications of the different types of data. There are also implications about PPP changes. Both 2005 and 2011 PPP are available when using PovcalNet.¹³ For South Asia and the year 2002, the official poverty line of \$1.25 and 2005 PPP provides a poverty rate of 44.10 percent and 637.87 million poor people. With the new \$1.90 and 2011 PPP, the poverty rate drops to 38.45 percent and 552.35 million poor people for the same year 2002, which makes a large difference. Quoting Deaton (2010), "PPP comparisons between widely different countries rest on weak theoretical and empirical foundations." So why not use simple exchange rates in U.S. dollars?

REFERENCES

- Atkinson, A. B., *The Economics of Inequality*, Clarendon Press, Oxford, 1983.
- Atkinson, A. B. and F. Bourguignon, "Poverty and Inclusion from a World Perspective," in J. E. Stiglitz and P.-A. Muet (eds), *Governance, Equity and Global Markets*, Oxford University Press, Oxford, 2001.
- Bauwens, L., M. Lubrano, and J.-F. Richard, *Bayesian Inference in Dynamic Econometric Models*, Oxford University Press, Oxford, 1999.
- Deaton, A., "Measuring Poverty in a Growing World (Or Measuring Growth in a Poor World)," *Review of Economics and Statistics*, 87, 1–19, 2005.
- , "Inequality, and the Measurement of World Poverty," *American Economic Review*, 100, 3–34, 2010.
- Deleeck, H., "The Adequacy of the Social Security System in Belgium, 1976–1985," *Journal of Social Policy*, 18, 91–117, 1989.
- Ferreira, F. H. G., S. Chen, A. Dabalen, Y. Dikhanov, N. Hamadeh, D. Jolliffe, A. Narayan, E. B. Prydz, A. R. P. Sangraula, U. Serajuddin, and N. Yoshida, "A Global Count of the Extreme Poor

¹³For 2005 PPP, see <http://iresearch.worldbank.org/PovcalNetPPP2005/index.htm>. From this website, it is possible to switch to 2011 PPP.

- in 2012: Data Issues, Methodology and Initial Results,” Policy Research Working Paper 7432, World Bank, Washington, DC, 2015.
- Flik, R. J. and B. M. van Praag, “Subjective Poverty Lines Definitions,” *De Economist*, 139, 311–30, 1991.
- Goedhart, T., V. Halberstadt, A. Kapteyn, and B. van Praag, “The Poverty Line: Concept and Measurement,” *Journal of Human Resources*, 12, 503–20, 1977.
- Greb, F., S. Klasen, S. Pasaribu, and M. Wiesenfarth, “Dollar a Day Re-Visited,” Discussion Paper 91, Georg-August-Universität, Göttingen, 2011.
- Hansen, B., “Sample Splitting and Threshold Estimation,” *Econometrica*, 68, 575–603, 2000.
- Holzmann, H., S. Vollmer, and J. Weisbrod, “Perspectives on the World Income Distribution—Beyond Twin Peaks Towards Welfare Conclusions,” Technical Report 158, Ibero America Institute for Economic Research (IAI), Göttingen, 2007.
- Kapteyn, A., P. Kooreman, and R. Willemse, “Some Methodological Issues in the Implementation of Subjective Poverty Definitions,” *Journal of Human Resources*, 23, 222–42, 1988.
- Pinkovskiy, M. and X. Sala-i-Martin, “Parametric Estimations of the World Distribution of Income,” Working Paper 15433, National Bureau of Economic Research, Cambridge, MA, 2009.
- Pradhan, M. and M. Ravallion, “Measuring Poverty Using Qualitative Perceptions of Consumption Adequacy,” *Review of Economics and Statistics*, 82, 462–71, 2000.
- Ravallion, M., *Poverty Comparisons*, Harwood Academic Press, Chur, Switzerland, 1994.
- Ravallion, M. and S. Chen, “How Did the World’s Poorest Fare in the 1990s?” *Review of Income and Wealth*, 47, 283–300, 2001.
- , “How Did the World’s Poorest Fare since the Early 1980s?” *World Bank Research Observer*, 19, 141–70, 2004.
- Ravallion, M., S. Chen, and P. Sangraula, “Dollar a Day Revisited,” *World Bank Economic Review*, 23, 163–84, 2009.
- Ravallion, M., G. Datt, and D. van de Walle, “Quantifying Absolute Poverty in the Developing World,” *Review of Income and Wealth*, 37, 345–61, 1991.
- Sala-i-Martin, X., “The World Distribution of Income: Falling Poverty and . . . Convergence, Period,” *Quarterly Journal of Economics*, 121, 351–97, 2006.
- Tendulkar, S. D., “Report of the Expert Group to Review the Methodology for Estimation of Poverty,” Technical Report, Government of India Planning Commission, New Delhi, 2009.
- Van den Bosch, K., T. Callan, J. Estivill, P. Hausman, B. Jeandier, R. Muffeis, and J. Yfantopoulos, “A Comparison of Poverty in Seven European Countries and Regions Using Subjective and Relative Measures,” *Journal of Population Economics*, 6, 235–59, 1993.
- World Bank, “Global Purchasing Power Parities and Real Expenditure 2005,” Technical Report, World Bank, Washington, DC, 2008.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher’s web-site:

- A. Simulation of a Bivariate Density Using a Grid