The Bayesian approach to poverty measurement

Lecture 2: Revising the IPL

Michel Lubrano

November 2022

Contents

1 Introduction

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. It is relatively easy to define a poverty line at the national level. We are here interested to define a poverty line at the international level. This is an important task for international institutions. For instance the United Nations have defined the project *Objectives for the Millennium*, whose aim was to cut extreme poverty in the world by half for 2015. It is far more difficult to define a poverty line at the world level. This is the object of this chapter which build on the work of the World Bank around the famous one dollar a day. The main difference between a national and an international poverty line (IPL) is that the former is devised for a single country when the latter should concern a group of countries, usually poor countries.

2 The concept of a poverty line

There are several ways of defining a poverty line and basically three.

An absolute poverty line is determined by reference to the cost of a given basket of goods or a minimum level of calories. 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 labour market.

A third concept is sometimes used which represents an intermediate case between these two polar cases: *A subjective poverty line* has 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 or Goedhart et al., 1977), 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 summarised as basic needs on one side and social inclusion for a given country on the other side (following Atkinson and Bourguignon, 2001).

3 A revised common poverty line for less developed countries

We now introduce the specificity of a poverty line for a group of countries.

3.1 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 wellbeing varies from country to country, simply because living standards, traditions, habits and other social characteristics are different (knowing that the PPP does not perfectly equalise 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, Chap. 10 or 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 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).

3.2 The econometric model of the World Bank

In a more recent paper, Ravallion et al. (2009) clearly identify two groups of countries in a new data set covering 74 developing countries with data collected over the period 1988-2005. They estimate a nonlinear 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:

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

$$s_i = \mathbb{1}(C_i < \theta) = 1 \text{ if } C_i < \theta, \text{ and 0 otherwise.} \quad (2)$$

s_i is an indicator function. For less developed countries, the elasticity of the poverty line with respect to mean consumption is assumed to be zero, that is we **impose** $\gamma_1 = 0$. 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, Ravallion et al. (2009) **impose** also a fixed $\theta = \$60$. With these two restrictions, they found that the revised IPL rises to $1.25 per day at 2005 PPP and to $1.90 using the new 2011 PPP.

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.

The data base used in Ravallion et al. (2009) is reported in the appendix of their paper. It concerns 74 developing countries. The data set includes national official poverty lines (PL) (or poverty lines computed by academics in some cases) and Private Consumption Expenditure (PCE) per capita collected in different years from 1988 to 2005. They have been adjusted by the household consumption PPP collected during the International Comparison Program of 2005 (World Bank, 2008). In Figure 1, we have plotted the

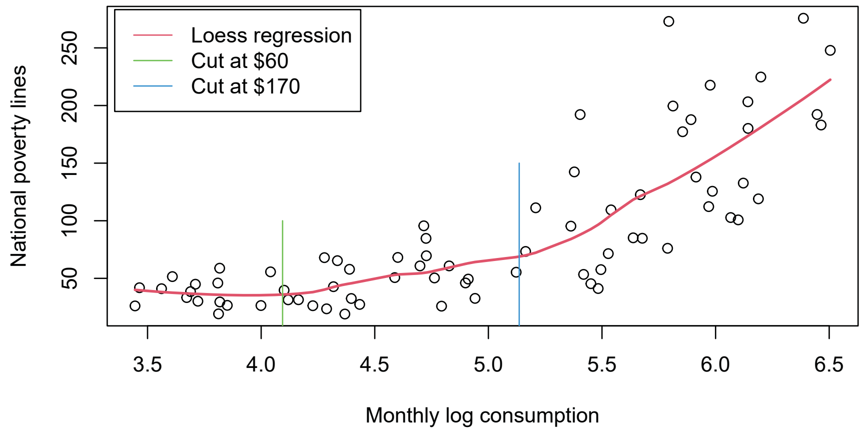

Figure 1: Ravallion et al. (2009) data set for the 74 developing countries

data used in Ravallion et al. (2009). This graph reproduces the plot given in their paper and representing the national poverty lines as a function of the log of national consumption per person at 2005 PPP. We have indicated on this plot two possible levels where an absolute poverty line should stop and consequently a relative poverty line should start. The non-parametric estimator chosen is the local polynomial fitting (`loess` command in `R`) with a smoothing parameter equal to 0.5.

3.3 Explaining differences in national poverty lines

The figures reported in the data base 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% of the mean consumption level, while it falls to 45% for the richer group of countries of the data base. This last figure is much more in accordance with the usual definition of a relative poverty line, usually half the mean income in Europe (or 60% of the median income). We arrive at the first figure of 92% 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 respectively $19 and $59, which means roughly between $0.60 and $2.00 a day.

How can we explain such differences? Deaton (2010) put forward the role of PPP calculations that we shall not detail here. In Xun and Lubrano (2018), we argue that national agencies may be influenced by arguments which are partly subjective and related to social inclusion when fixing the national poverty line. 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. This militates for a less restrictive specification of the model, which would then relate the national poverty lines to a set of explanatory variables including mean consumption and eventually the rate of unemployment as a possible indicator for social inclusion.

Even for an absolute poverty line based on a minimum number calories, 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. Social inclusion is recognised as an important factor, 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*".

3.4 A model for an international "subjective" poverty line

The complete and extended model that we consider is:

$$
\begin{aligned}
z_i &= s_i(\alpha_1 + \gamma_1 \log C_i + \beta_1 x_i) + (1 - s_i)(\alpha_2 + \gamma_2 \log C_i) + \epsilon_i \quad (3)\\
s_i &= \begin{cases} 1 & \text{if } C_i < \theta \\ 0 & \text{otherwise} \end{cases} \quad (4)\\
\text{Var}(\epsilon_i) &= s_i \sigma_1^2 + (1 - s_i)\sigma_2^2 \quad (5)
\end{aligned}
$$

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 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:

$$\mathrm{E}(z_i|s_i = 1) = \alpha_1 + \gamma_1 \mathrm{E}(\log C_i|s_i = 1) + \beta_1 \mathrm{E}(x_i|s_i = 1). \tag{6}$$

In words, the poverty line 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.

1. 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).

2. Second, as it depends both on consumption level and 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.

3. Third, our poverty line is defined with respect to a common group that each country is supposed to identify with. 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 equalised using either 2005 PPP or 2011 PPP.

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 from the reference group (like China for instance) 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 $\log C_i$.

4 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, Chap. 8), namely:

$$\begin{aligned} \mathrm{E}(z_i|x_i) &= x_i'\beta_1 \quad \text{if } C_i \leq \theta, \\ \mathrm{E}(z_i|x_i) &= x_i'\beta_2 \quad \text{if } C_i > \theta. \end{aligned}$$

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:

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

Regrouping these elements in a single equation, we get:

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

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:

$$X(\theta) = [s_i x_i', (1 - s_i)x_i'], \tag{7}$$

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

$$z = X(\theta)\beta + \epsilon, \tag{8}$$

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

4.1 Likelihood and posteriors

Considering N observations, the likelihood function of model (8) is:

$$L(\beta, \sigma^2, \theta; z) \propto \sigma^{-N} \exp\left[-\frac{1}{2\sigma^2}\sum_{i=1}^{N}[z_i - X_i'(\theta)\beta]^2\right]. \tag{9}$$

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:

$$\begin{aligned}\pi(\beta|\sigma^2) &= f_N(\beta_0, \sigma^2 M_0^{-1}),\\ \pi(\sigma^2) &= f_{Ig}(\sigma^2|\nu_0, s_0).\end{aligned}$$

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

$$\pi(\beta|\theta, z) = f_t(\beta|\beta_*(\theta), s_*(\theta), M_*(\theta), \nu_*), \tag{10}$$

$$\pi(\sigma^2|\theta, z) = f_{Ig}(\sigma^2|\nu_*, s_*(\theta)), \tag{11}$$

$f_t(.)$ being the Student distribution. The different posterior hyperparameters are defined by:

$$\begin{aligned}M_*(\theta) &= M_0 + X'(\theta)X(\theta),\\ \beta_*(\theta) &= M_*^{-1}(\theta)[X'(\theta)z + M_0\beta_0],\\ s_*(\theta) &= s_0 + \beta_o' M_0\beta_0 + z'z - \beta_*'(\theta)M_*(\theta)\beta_*(\theta),\\ \nu_* &= \nu_0 + N.\end{aligned}$$

The usual and convenient practice is to use a noninformative prior for the regression parameters with $\beta_0 = 0$, $M_0 = 0$, $s_0 = 0$, $\nu_0 = 0$. 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 (10), being simply proportional to the inverse of its integration constant times the prior density of θ:

$$\pi(\theta|z) \propto |s_*(\theta)|^{-(N-k)/2}|M_*(\theta)|^{-1/2}\pi(\theta). \tag{12}$$

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:

$$\pi(\beta|z) = \int f_t(\beta|\beta_*(\theta), s_*(\theta), M_*(\theta), \nu_*)\, \pi(\theta|z)\, d\theta, \tag{13}$$

and

$$\pi(\sigma^2|z) = \int f_{Ig}(\sigma^2|\nu_*, s_*(\theta))\, \pi(\theta|z)\, d\theta. \tag{14}$$

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.[1] As (12) 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, noted $\theta^{(j)}$ for $j = 1, \dots, M$. Briefly, the grid over which to evaluate (12) 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 (13), which is a Student density.

4.2 The two variance case

For modelling 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:

$$\mathrm{Var}(\epsilon_i) = s_i \sigma_1^2 + (1 - s_i)\sigma_2^2. \tag{15}$$

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

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

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

$$\begin{aligned} z(\theta, \phi) &= [z_i/\sqrt{h(\theta, \phi)}], & (16)\\ X(\theta, \phi) &= [s_i x_i'/\sqrt{h(\theta, \phi)}, (1 - s_i)x_{it}'/\sqrt{h(\theta, \phi)}]. & (17) \end{aligned}$$

[1] It 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. Using deterministic integration rules leads to much more complicated procedures.

The regression model becomes:

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

its likelihood function being

$$\begin{aligned} L(\beta, \sigma^2, \theta, \phi; z) \quad &\propto \quad \sigma^{-N} \prod_{i=1}^{N} h_i(\theta, \phi)^{-1/2} \times \\ & \exp\left[-\frac{1}{2\sigma^2} (z(\theta, \phi) - X(\theta, \phi)\beta)'(z(\theta, \phi) - X(\theta, \phi)\beta) \right]. \end{aligned} \tag{18}$$

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:

$$\pi(\theta, \phi|z) \propto \prod_{i=1}^{N} h_i(\theta, \phi)^{-1/2} |s_*(\theta, \phi)|^{-(N-k)/2} |M_*(\theta, \phi)|^{-1/2} \pi(\theta)\pi(\phi). \tag{19}$$

It is slightly more difficult to draw $\theta^{(j)}$ and $\phi^{(j)}$ jointly from this bivariate density (19) than to draw $\theta^{(j)}$ from the univariate density (12). A feasible method can be found if we remember that it is always possible to decompose a bivariate density into:

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

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 next section.

4.3 Simulation of a bivariate density using a grid

Let us consider the bivariate posterior density:

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

We know the analytical form of the joint density $\pi(\phi, \theta|z)$, but neither its marginal $\pi(\theta|z)$ nor its conditional $\pi(\phi|\theta, z)$. We want to draw random numbers for the joint posterior density. To do so, we first evaluate this bivariate density on a grid, filling a matrix F where the rows correspond to θ and the columns to ϕ. From this matrix of points, we can numerically determine the marginal density $\pi(\theta|z)$ by summing over the columns. Using this marginal density and using the inverse transformation method, we can draw a value for θ. For a given draw of θ, we have to find the corresponding conditional density $\pi(\phi|\theta, z)$ as a row of matrix F. Of course, the draw will not correspond exactly to one of the predetermined points of the grid in θ. So we have to proceed by linear interpolation between two lines.

1. Compute the cumulative numerically and then use the inverse transformation method to draw $\theta^{(j)}$ from $\pi(\theta|z)$.

2. Find the two nearest points of $\theta^{(j)}$ on the grid of θ, denoted as $\theta^{(j-)}$ and $\theta^{(j+)}$.

3. Calculate the differences: $a = \theta^{(j)} - \theta^{(j-)}$, $b = \theta^{(j+)} - \theta^{(j)}$ and $c = |\theta^{(j+)} - \theta^{(j-)}|$.

4. Obtain the conditional posterior densities $\pi(\phi|\theta^{(j-)}, z)$ and $\pi(\phi|\theta^{(j+)}, z)$ from the joint posterior matrix F.

5. Compare each point of the two above conditional posterior densities in order to get $\pi(\phi|\theta^{(j)}, z)$ by linear interpolation:

$$\pi(\phi|\theta^{(j)}, y) = \begin{cases} \pi(\phi_k|\theta^{(j-)}, z) + a \times (\pi(\phi_k|\theta^{(j+)}, z) - \pi(\phi_k|\theta^{(j+)}, z))/c & \text{if } \pi(\phi_k|\theta^{(j+)}, z) \geq \pi(\phi_k|\theta^{(j-)}, z), \\ \pi(\phi_k|\theta^{(j+)}, z) + b \times (\pi(\phi_k|\theta^{(j-)}, z) - \pi(\phi_k|\theta^{(j+)}, z))/c & \text{otherwise,} \end{cases}$$

   and impose normalisation of this conditional density:

$$\sum_{k=1}^{k} \pi(\phi_k|\theta^{(j)}, z) = 1,$$

   assuming that ϕ_k is the k^{th} point on the grid of ϕ.

6. Compute the cumulative numerically and then use the inverse transformation method to draw $\phi^{(j)}$ from $\pi(\phi|\theta^{(j)}, z)$.

7. Store the jth joint draw : $(\theta^{(j)}, \phi^{(j)})$.

4.4 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

$$\mathrm{E}(z_i|s_i = 1) = \gamma_1 \mathrm{E}(\log C_i|C_i < \theta) + \beta_1 \mathrm{E}(x_i|C_i < \theta). \tag{20}$$

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

$$\mathrm{E}(z_i|s_i = 1) = \gamma_1 \mathrm{E}(w_i \times \log C_i|C_i < \theta) + \beta_1 \mathrm{E}(w_i \times x_i|C_i < \theta). \tag{21}$$

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 (19). 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 get 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 $\sigma^{(j)}$:

$$z^{(j)} = \gamma_1^{(j)} \sum_{i=1}^{n} w_i \log(C_i) \mathbb{1}(C_i < \theta^{(j)}) + \beta_1^{(j)} \sum_{i=1}^{n} w_i x_i \mathbb{1}(C_i < \theta^{(j)}), \tag{22}$$

where w_i are weights (equal or unequal) summing to one according to the scheme $\sum_i w_i \times \mathbb{1}(C_i < \theta^{(j)}) = 1$. 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.

5 Empirical comparison of different Bayesian poverty lines

We now present our empirical results for three different cases. First, 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).

5.1 The two regime model

The most general model we start with is:

$$\begin{aligned} z_i &= s_i(\alpha_1 + \gamma_1 \log C_i + \beta_1 Ur_i) \\ &\quad +(1 - s_i)(\alpha_2 + \gamma_2 \log C_i + \beta_2 Ur_i) + \epsilon_i, \end{aligned} \tag{23}$$

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.[2] 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 nonlinear, first because the value $\phi = 1$ does not belong to a 90% 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,

[2] Unemployment as a percentage of the total labour force, probably only covering the formal sector. The source is the World Bank web site.

Table 1: Explaining national poverty lines using
a two regime model with two variances, 2005 PPP

	Unweighted		Population weighted		Nber Poor weighted	
	Mean	s.d.	Mean	s.d.	Mean	s.d.
			First regime			
γ_1	8.64	(0.72)	10.26	(0.59)	10.30	(0.61)
Ur	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 bottom panel indicates the average weights given to China and India in the three different weighting schemes.

90% 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 are really needed.

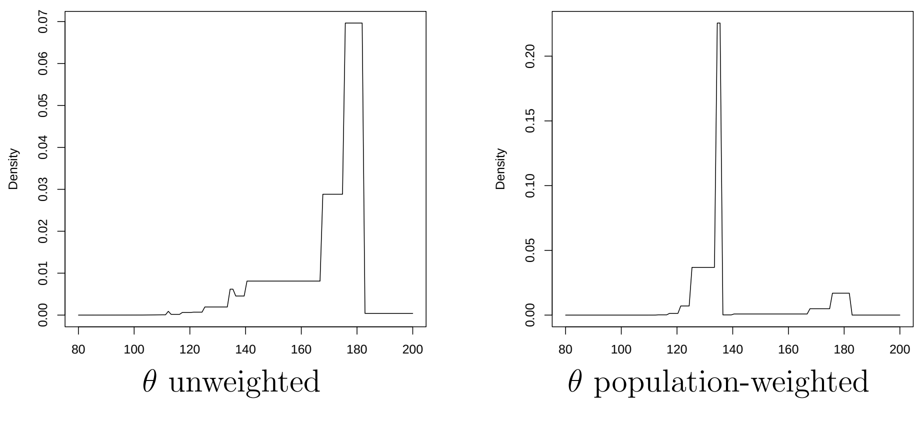

Figure 2: Posterior density of θ for the unweighted and weighted cases

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% 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% 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 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

Table 2: 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	34	2	33
weighted	Bangladesh	China	Brazil
	India		Mexico
	Pakistan		Russia
Nber poor	36	3	32
weighted	Bangladesh	Indonesia	Brazil
	China		Mexico
	India		Russia
	Pakistan		

Are listed countries with more than 100M inhabitants. Numbers indicate the average size of the group.

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 modelled international poverty line.

5.2 Modelling the poverty line

What is the real influence of population weighting when determining the IPL? If we simply compute the mean of 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 or a population weighted mean. In the first case, we get $1.48 (0.020) while in the other case we get $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.

Let us now report our inference results concerning our modelled IPL. Equations (20) and (21) define the poverty line as a linear function of the

Table 3: Which International Poverty Line?

Regression model	Modelled IPL, 2005 PPP	*Modelled IPL, 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 modelled IPL for each model in column one, using 2005 PPP. Draws for the modelled IPL are obtained using (22). The last column in italics corresponds to 2011 PPP and is reported only as an indication.

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% 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% 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% 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 modelled 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, so again a greater value than \$1.90 IPL of the World Bank. But these figures have to be taken with a grain of salt in accordance with the remarks that are usually made concerning the 2011 PPP conversion, including those of the World Bank itself. See Xun and Lubrano (2018) for more details.

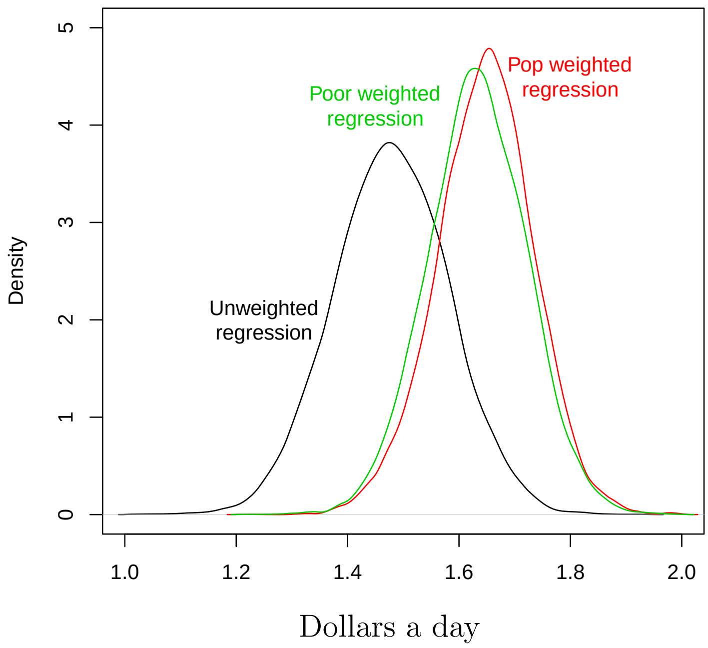

Figure 3: Posterior density of the modelled IPL using 2005 PPP

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 depending 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 like 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 have followed his advice.

6 Conclusion

Using the same data as Ravallion et al. (2009), we have provided 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 anyway greater than \$1.90. Whatever the PPP, the World Bank always underestimate the IPL because it neglects social inclusion. Another finding is that weighting strongly affects the final result, and that the change depends greatly on the model specification for deriving the IPL. 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.

A Bayesian approach has forced us to explicit a certain number of assumptions and has shown the consequences of these assumptions.

References

Atkinson, A. B. (1983). *The Economics of Inequality*. Clarendon Press, Oxford.

Atkinson, A. B. and Bourguignon, F. (2001). Poverty and inclusion from a world perspective. In Stiglitz, J. E. and Muet, P.-A., editors, *Governance, Equity and Global Markets*. Oxford University Press.

Bauwens, L., Lubrano, M., and Richard, J.-F. (1999). *Bayesian Inference in Dynamic Econometric Models*. Oxford University Press, Oxford.

Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics*, 87:1–19.

Deaton, A. (2010). Inequality, and the measurement of world poverty. *American Economic Review*, 100:3–34.

Goedhart, T., Halberstadt, V., Kapteyn, A., and Van Praag, B. (1977). The poverty line: concept and measurement. *Journal of Human Resources*, 12(4):503–520.

Kapteyn, A., Kooreman, P., and Willemse, R. (1988). Some methodological issues in the implementation of subjective poverty definitions. *Journal of Human Resources*, 23:222–242.

Pradhan, M. and Ravallion, M. (2000). Measuring poverty using qualitative perceptions of consumption adequacy. *Review of Economics and Statistics*, 82:462–471.

Ravallion, M. and Chen, S. (2001). How did the world's poorest fare in the 1990s. *Review of Income and Wealth*, 47:283–300.

Ravallion, M. and Chen, S. (2004). How did theworld's poorest fare since the early 1980s. *World Bank Research Observer*, 19:141–170.

Ravallion, M., Chen, S., and Sangraula, P. (2009). Dollar a day revisited. *World Bank Economic Review*, 23:163–184.

Ravallion, M., Datt, G., and van de Walle, D. (1991). Quantifying absolute poverty in the developing world. *Review of Income and Wealth*, 37:345–361.

Tendulkar, S. D. (2009). Report of the expert group to review the methodology for estimation of poverty. Technical report, Government of India Planning Commission, New Dehli.

Van den Bosch, K., Callan, T., Estivill, J., Hausman, P., Jeandidier, B., Muffeis, R., and Yfantopoulos, J. (1993). A comparison of poverty in seven European countries and regions using subjective and relative measures. *Journal of Population Economics*, 6:235–259.

World Bank (2008). Global purchasing power parities and real expenditure 2005. Technical report, World Bank, Washington DC.

Xun, Z. and Lubrano, M. (2018). A Bayesian measure of poverty in the developing world. *Review of Income and Wealth*, 64(3):649–678.