

Are the World's Poorest Being Left Behind?

Martin Ravallion¹

Department of Economics, Georgetown University, and NBER
Washington DC., 20057, U.S.A.
mr1185@georgetown.edu

Abstract: Traditional assessments of economic growth and progress against poverty tell us little about whether the poorest are being left behind—whether the consumption floor is rising above the biological minimum. To address this deficiency, the paper identifies the expected value of the floor as a weighted mean of observed consumptions for the poorest stratum. Under the identifying assumptions and using data for the developing world over 1981-2011, the estimated floor is about half the \$1.25 a day poverty line. Economic growth and social policies have delivered only modest progress in raising the floor, despite overall growth and progress in reducing the number living near the floor.

Keywords: Poverty; consumption floor; Rawls; growth; safety-nets

JEL: I32, I38, O15

¹ For helpful discussions on the topic of this paper and comments on the paper the author is grateful to Francois Bourguignon, Mary Ann Bronson, Cait Brown, Shaohua Chen, Denis Cogneau, James Foster, Garance Genicot, Peter Lanjouw, Nkunde Mwase, Henry Richardson, Dominique van de Walle and seminar participants at the International Monetary Fund, the Paris School of Economics, the International Labour Organization, the University of Antwerp, the Human Development and Capabilities Association Conference 2015, Georgetown University, and the 30th annual conference of WIDER, Helsinki, 2015. The author also thanks the journal's editor and five referees for their comments. One of the referees was especially helpful in developing a stronger rationale for the proposed estimator in Section 4.

1. Introduction

At the launch of the United Nations' (2011) [Millennium Goals Report](#), the U.N.'s Secretary-General Ban Ki-moon said that:

“The poorest of the world are being left behind. We need to reach out and lift them into our lifeboat.”

This view is heard often. A press release by the [International Food Policy Research Institute](#) carried the headline: “The world’s poorest people not being reached.”² On the occasion of the International Day for the Eradication of Poverty in 2014, the [International Labor Organization](#)'s Director-General, Guy Ryder, wrote that “Poverty is not yet defeated. Far too many are being left behind.” And the Vatican’s representative to the United Nations reaffirmed in 2015 that the poorest of the world are being left behind.³

Yet other observers appear to tell a very different story. They use aphorisms such as “a rising tide lifts all boats” or that “growth is good for the poor” (the title of a paper published in this journal) or that the poor are “breaking through from the bottom.”⁴ Proponents of this alternative view draw on data in the form of survey-based measures of poverty or inequality.

This paper tries to understand these seemingly conflicting views. The central issue is how we should assess progress against poverty. The main approach of economists and statisticians has been to count the poor in some way. One might track the proportion of the population living below some deliberately low poverty line or use a more sophisticated measure giving higher weight to poorer people. Bowley (1915) was a prominent early advocate of this approach. Since then, the theoretical foundations of the approach are found in a large literature on poverty measurement, in which various axioms have been proposed.⁵ I dub this the counting approach. The common feature is that the measure of poverty is obtained by counting poor people, though

² The press release was for an IFPRI report Ahmed et al. (2007).

³ Quoted by James (2015).

⁴ The first expression is attributed to John F. Kennedy, the middle claim is the title of an influential paper by Dollar and Kraay (2002), reiterated by Dollar et al. (2013), while the last expression is due to Radelet (2015).

⁵ The most commonly used axioms are: (i) focus: the measure of poverty should be unaffected by any changes in the incomes (or consumptions) of those who are not deemed to be poor ; (ii) monotonicity: holding all else constant, the measure of poverty must rise if a poor person experiences a drop in her income; (iii) subgroup monotonicity: poverty rises when any sub-group becomes poorer; (iv) scale invariance: the measure is unchanged when all incomes and the poverty line increase by the same proportion; (v) the transfer principle: the measure of poverty falls whenever a given sum of money is transferred from a poor person to someone even poorer. An influential early contribution to the axiomatic foundations was made by Sen (1976), although Sen’s measure did not satisfy all of the above axioms. Other axioms have also been proposed; for a fuller listing see Foster et al. (2013).

possibly giving higher weight to the poorest. Typically, the measure of poverty is a population-weighted average of individual measures of poverty across the population (counting the non-poor as having zero poverty).

The paper argues that the counting approach does not adequately address prevailing concerns about whether the poorest are being left behind during economic growth. Logically, for the poorest to not be left behind there must be an increase in the lower bound of the support of the distribution of levels of living. That lower bound can be called the consumption floor. An appealing concept of “level of living” in this context is permanent consumption (Friedman, 1957). If the poorest person sees a gain in permanent consumption then (by definition) the consumption floor has risen. Human physiology makes it plausible that there is a positive floor, given the nutritional requirements for basal metabolism and normal activities. This can be called the “the biological floor.” However, given economic growth and (private and public) redistribution the actual consumption floor may well be above the biological floor.

Prevailing measures of economic progress put little or no weight on progress in raising the floor. It is clear that the growth rate in the overall mean consumption will have a low (possibly very low) implicit weight on the growth rate in the floor, given that the share going to the poorest is likely to be low. What about measures of poverty? Each panel of Figure 1 gives two cumulative distribution functions (CDFs). In each case, the upper CDF is the initial one and the lower CDF is for a later date. As drawn, no standard measure of poverty will show an increase (Atkinson, 1987). The drop in the incidence of poverty at the indicated poverty line is similar in panels (a) and (b). In (a), the counting approach can reasonably claim that many of the poorest have been reached even though the floor has not risen, so some people still remain living at the same very low level. In panel (b), the same reduction in the poverty rate has come with a rising floor—implying that the poorest are left behind.

The idea that we should judge progress in part by success in raising the floor is missing from all standard poverty measures. The concept of the consumption floor is conceptually distinct from existing poverty lines.⁶ Naturally, any poverty line aims to reflect what “poverty” means in a specific society, on the understanding that (potentially many) people live below that level. The poverty line is a normative concept, while the consumption floor is a positive one. The most widely-used poverty measure, the “headcount index” attaches no value to success in raising

⁶ For further discussion of poverty lines in theory and practice see Ravallion (2012).

the floor.⁷ Those living at the floor will have the highest weight in the subset of poverty measures that penalize inequality among the poor though even then there is no assurance that adequate weight will be attached to progress in raising the floor.⁸ Indeed, in most measures the weight on the individual poverty measure is the corresponding population density, which may well be quite low for the poorest.

This neglect of explicit attention to the poorest appears to be due, at least in part, to the difficulties in identifying the floor.⁹ While some theoretical formulations of the policy-evaluation problem have assumed that measuring the lowest level of living is straightforward,¹⁰ that is clearly not the case in practice. The lowest observed consumption in a survey may differ appreciably from the lowest permanent consumption. Given the current interest in assuring that no one is left behind, this is a gap in the “dashboard” of development indicators.

The difference between the two approaches, as illustrated by Figure 1, begs some questions for which we currently have little idea of the answers: The consumption floor plausibly exists, but at what level? Has the growth we have seen in mean consumption and income in the developing world come with growth in the level of the floor? Has success against poverty judged by the counting approach also come with success in raising the floor?

The task of addressing these questions calls for a method of estimating the level of the consumption floor. With ideal data—let us imagine a complete set of accurate measures of permanent consumption across the population—the problem is trivial; we just look for the lowest value. However, actual data are a very long way from this ideal. Instead, we confront a severe and fundamental data constraint. Standard household surveys only ask respondents about consumption or income over relatively short recall periods, and such data are certain to contain sizable transient components. It is also clear that poverty monitoring and social policy discussions are motivated by concerns about low levels of typical consumption. When we refer to the typical level of living of the poorest stratum we are acknowledging that consumption may be low at one date for transient reasons. Identifying the floor as the strict lower bound of

⁷ This reflects well known limitations of this measure, which fails both the monotonicity and transfer axioms. The income share of the poorest $x\%$ has been used as a measure of inequality, but it is also known to have deficiencies when judged by the standard axioms of inequality measurement; for further discussion see Fields (2001, Chapter 2).

⁸ The limit of the Foster-Greer-Thorbecke (1984) measure as their inequality aversion parameter goes to infinity is the lowest value level in the data. This is only the floor, as measured here, if one is certain that the lowest observed value is the lower bound to permanent consumption. It is argued below that this is questionable.

⁹ See, for example, Freiman’s (2012) comments on Rawls’s difference principle.

¹⁰ See, for example, the discussion in Fleurbaey and Maniquet (2011, Chapter 12).

observed consumptions or incomes would clearly be unsatisfactory as it is subject to idiosyncratic transient factors, and possibly sizeable measurement errors. We need a more robust approach that is still operational with the data available.

The paper proposes an approach that can be implemented with readily available data. The method entails estimating the expected value of the consumption floor as a weighted mean formed over a stratum of people with low observed consumption levels, with higher weight for those who appear to be the poorer. These are not normative weights, but rather they reflect the likelihood that any observed consumption is in fact the floor. Specifically, under weak assumptions, it is shown that the lowest observed consumption has the highest probability of being at the floor, but that probability is less than one. When parameterized by a power function, the idea of the consumption floor can then be interpreted in terms of standard, readily available, poverty measures in the Foster, Greer and Thorbecke (1984) class. For various parameter values, the implied estimate of the consumption floor is around half of the international poverty line of \$1.25 a day. The paper also finds that this estimate agrees quite closely with alternative sources, including an approach based on national poverty lines, when interpreted as the expected value of the consumption floor plus a relative component proportional to mean consumption.

The paper then shows that, while the developing world has seen a high growth rate in mean consumption in the new millennium, and the counting approach shows much progress in reducing poverty, there has been little progress in raising the floor over the last 30 years. The distribution of the gains amongst the poor has meant that the expected value of the lowest level of living has advanced rather little.

After reviewing the literature and policy discussions related to the idea of a consumption floor (Section 2), the paper describes the data to be used in this study (Section 3). Then it turns to the proposed measure of the floor and its empirical implementation (Sections 4) as well as discussing the alternative approaches (Section 5). For comparison purposes, the paper then presents new evidence using the counting approach (section 6). In the light of the paper's main findings, Section 7 offers some observations on the coverage of social safety nets in developing countries. Section 8 concludes.

2. The consumption floor in theory and policy

Moral philosophers have long argued that justice is only served when every individual is covered by its precepts—none are left behind. An application to distributive justice assesses a society’s economic progress by its ability to enhance the economic welfare of the least advantaged, following the two principles of justice proposed by Rawls (1971). First, each person should have an equal right to the most extensive set of liberties compatible with the same rights for all. Second, subject to that constraint, social choices should only permit inequality if it is efficient to do so—that a difference is only allowed if both parties are better off as a result; this is what Rawls called the “difference principle.” By this view, a higher floor (as in Figure 1(b)) is not only preferred, it is the criterion by which we judge progress.¹¹ Rawls’s difference principle is often interpreted as “maximin”—to maximize the minimum level of welfare. However, Rawls recognized that this is almost certainly unworkable in practice, as it is too demanding to know who is literally the poorest. Rawls does not appear to have imagined that household survey data could be used, but his concerns would also apply to such data.

Rawls (1971) claimed that his difference principle would be agreed among equals in a “veil of ignorance” about where they would find themselves in the real world.¹² This claim stimulated much debate. Harsanyi (1975) questioned whether maximin was a more plausible choice for a social contract than maximizing average utility even behind the veil of ignorance unless there was extreme risk aversion. Roemer (1996, Chapter 5) also questioned whether maximin would emerge as the solution. These critiques rested on the assumption that agents behind the veil would maximize expected utility, which depends solely on their own consumption (and leisure). However, one can defend Rawls’s difference principle without accepting his rationale in terms of a social contract formed behind the veil of ignorance. Hammond (1976) showed that a generalized lexicographic version of maximin, dubbed leximin in the literature, can be derived from a set of axioms including a requirement that reducing the disparities in welfare between the rich and the poor is socially preferred, other things being equal. Similarly, Fleurbaey and Maniquet (2011, Chapter 3) show that leximin is implied by

¹¹ While popularity need not guide ethical judgments, it is at least notable in the context of understanding debates about distributive justice that there is experimental evidence indicating that a non-negligible number of people make distributional judgments consistently with a Rawlsian “maximin” criterion (Michelbach et al., 2003).

¹² The veil of ignorance was a thought device to assure that morally irrelevant—inherited or acquired—advantages in the real world did not color judgments about distributive justice.

what they call the “priority among equals” axiom. Again this requires that more equitable allocations are socially preferred but that (echoing Rawls) this never trumps efficiency in that a situation in which everyone is better off is always preferred. Roemer (2014) argued for leximin but from a somewhat different starting point, namely the desire to equalize opportunities.

The Rawlsian approach also has deep roots in development and social-policy thinking. The widespread concerns that a growing economy will leave the poorest behind have led to social policies that strived to support consumption levels above the biological minimum. Indeed, this has long been a guiding principle in rich and poor countries alike. One motivation for the laws establishing statutory minimum wage rates that first appeared in the late 19th century is that they raise the consumption floor.¹³ Social policies have often aimed to guarantee a minimum income.¹⁴ There have also been advocates of the idea of a “basic-income guarantee”—a fixed cash transfer to every adult person. The International Labor Organization (2012) has recommended a comprehensive “Social Protection Floor,” comprising “nationally defined sets of basic social security guarantees” spanning health, schooling and income security.¹⁵

In the new millennium, mainstream development policies have come to embrace a range of direct interventions, variously called “antipoverty programs,” “social safety nets,” and “social assistance;” here I call them social safety nets (SSN’s).¹⁶ Their common feature is the use of direct income transfers to poor families. While this was rare in the developing world prior to the mid-1990s, today almost every country has at least one SSN program (World Bank, 2014). The new SSN programs have mainly been in the form of conditional cash transfers and workfare schemes (World Bank, 2014). The compilation of survey-based estimates of SSN coverage spanning 2000-2010 in the World Bank’s [ASPIRE](#) database indicates that the share of the population receiving help from SSN programs is growing rapidly, although there are probably selection biases in the data. The term “safety net” evokes the idea of a floor. Indeed, raising the consumption floor is a common motivation for SSN programs; examples include the two largest

¹³ There are also well-known efficiency arguments, notably in non-competitive labor markets. The first minimum wage law was introduced by New Zealand in 1894.

¹⁴ An early example was the Speenhamland System of 1795 introduced by the justices of Berkshire, which guaranteed local working-class residents a basic income indexed to the price of bread; for further discussion of this and subsequent policies see Ravallion (2016, Chapter 10).

¹⁵ On the arguments for a social security floor see Cichon and Hagemeyer (2007).

¹⁶ A good working definition is: “Social safety nets are non-contributory transfers designed to provide regular and predictable support to targeted poor and vulnerable people.” (World Bank, 2014, p.xii.) Also see Barrientos (2013).

programs to date in population coverage, namely China's *Di Bao* program and India's *National Rural Employment Guarantee Scheme*.¹⁷

The fact that SSN coverage is expanding gives hope that the floor is rising. Of course, whether this is happening in practice is another matter. To assess whether we are seeing progress against poverty consistently with the Rawlsian approach, one needs to define and measure the consumption floor. No such definition and measure is currently available. It is widely acknowledged that there is a need to focus on more than the growth rate of the overall mean, and descriptive tools such as the growth incidence curve of Ravallion and Chen (2003) have clearly helped. However, there has been little effort to study the growth rate of the lower bound of the distribution of levels of living, even though that lower bound has received much attention in social policy discussions, social choice theory and moral philosophy.

While economists measuring poverty have not attached much weight to the level of the floor, the concept has played a role in positive economics. Early ideas in classical economics of the "subsistence wage" can be interpreted as the wage rate required to assure that the biological floor is reached for a typical family.¹⁸ Famously, Malthus (1806) argued that the economic dynamics of population growth assures that the unskilled wage rate stays at the subsistence level; any temporary increase (decrease) in the consumption of working-class families in a neighborhood of the floor would induce population growth (contraction). The idea of a floor has been a feature of development models since Lewis (1954).

The idea has continued to play a role in modern economics. It has been built into demand models, such as the widely-used linear expenditure system. The idea is found in modern theoretical treatments of the problem of determining the optimal population size.¹⁹ The idea of a consumption floor is also found in modern dynamic models.²⁰ For example, some theoretical models have postulated an instantaneous utility function of the Stone–Geary form; consumers then maximize the present value of the utility stream subject to their consumption not falling

¹⁷ The *Di Bao* program makes transfers to bring urban residents up to locally determined "*Di Bao* lines" (see, for example, Ravallion, 2014b). The *Rural Employment Guarantee Scheme* in India aims to guarantee up to 100 days of work per household per year doing unskilled manual labor at stipulated minimum wage rates; see Dutta et al. (2014). The latter program can be interpreted as an attempt to enforce the minimum wage rate in an informal economy.

¹⁸ See Blaug's (1962) discussion of the classical model of wage determination.

¹⁹ See Dasgupta (1993, Chapter 13). Blackorby and Donaldson (1984) proposed that social welfare increases with a larger population if and only if the extra people have a level of consumption above a critical minimum. This can be interpreted as an ethical floor, unlike the consumption floor, which is a positive concept.

²⁰ See, for example, Azariadis (1996), Ben-David (1998) and Kraay and Raddatz (2007).

below the floor (in addition to other standard constraints).²¹ There are also arguments on the production side, whereby the existence of a floor generates a low-level non-convexity in production possibility sets. Various theoretical arguments have been made along these lines. The essential idea is that worker productivity and/or access to credit (given default chances) suffer when a person's consumption is close to the floor.²²

Under certain conditions, low-level non-convexities (such as from thresholds) can also generate dynamics with multiple equilibria in which the low-level attractor among a set of steady-state equilibria is identified as a poverty trap.²³ This can explain why some people may be stuck at the floor; the dynamic adjustment in response to small transient gains will push them back to the floor in due course. It can also explain why limited coverage and low transfers can entail that SSN programs are ineffective in raising the floor. Only with a sufficiently large gain will the poorest be able to attain a growth path toward their preferred long-run equilibrium.

Given the prominence of the idea of a consumption floor in moral philosophy and social policy, as well as positive economics, it is of interest to see how one might make the idea operational—to quantify the expected level of the floor and how it has evolved over time. That is the task of the rest of the paper.

3. The data

The paper focuses mainly on the developing world over the period 1981-2011. The primary data source is the World Bank's [PovcalNet](#) website. Here only a brief summary is provided.²⁴

The database of *PovcalNet* draws on 900 surveys spanning 125 developing countries. Using the most recent survey for each country, 2.1 million households were interviewed. The surveys are not of course evenly spaced over time, and they do not span the full 30 years period for all countries. The surveys were mostly done by governmental statistics offices as part of their routine operations. Not all available surveys are included in *PovcalNet*. A survey was dropped if there were known to be serious comparability problems with the rest of the data set. Obvious

²¹ An example is found in Lopez and Servén (2009), who add a subsistence consumption parameter to the type of model discussed in Aghion et al. (1999).

²² Examples include Mirrlees (1975), Stiglitz (1976), Dasgupta and Ray (1986), Lipton (1988) and Banerjee and Newman (1994).

²³ See, for example, Galor and Zeira (1993). A review of alternative model can be found in Azariadis (2006).

²⁴ The sources and estimation methods are described in greater detail in Chen and Ravallion (2010).

problems were addressed by either re-estimating the consumption/income aggregates or by dropping a survey. Of course, there are data problems that cannot be dealt with, and differences in survey methods can create differences in the estimates obtained.

All poverty measures are estimated from the primary (unit record or tabulated) sample survey data rather than relying on pre-existing estimates. Prior truncations of the data (trimming the bottom or top) are avoided as far as possible, and appear to be rare at the bottom of the distribution. Past estimates are updated to ensure internal consistency with new data.²⁵ All distributions are weighted by household size and sample weights.

Households are ranked by either consumption or income per person, with consumption being preferred. About 70% of the surveys allow a consumption-based measure. Estimates are also done excluding the income surveys. The measures of consumption (or income, when consumption is unavailable) are reasonably comprehensive, including both cash spending and imputed values for consumption from own production. However, as in virtually all household surveys, the recall periods for consumption in the survey questionnaire are short. Food, for example, is typically asked for the week or two prior to the interview. And if it is an income survey, then there is no strict lower bound to observed income in a relatively short time period. (Zero income is common, but income in the survey period can also be negative.)

The poverty count is the number of people living in households with per capita consumption below the international poverty line. All currency conversions are at purchasing power parities using the results of the 2005 round of the International Comparison Program. The main international poverty line is \$1.25 a day as proposed by Ravallion et al. (2009) who provide various rationales for this line.

In addition to the popular headcount index (count of poor people normalized by the population), *PovcalNet* provides the poverty gap index and the squared poverty gap index. All three of these measures are members of the class of measures introduced by Foster, Greer and

Thorbecke (FGT) (1984), which can be written as $P_\alpha = \sum_{y_i \leq z} (1 - y_i / z)^\alpha / n$ where y_i is

consumption (or income) of person $i=1, \dots, n$, z is the poverty line, and α is a non-negative parameter. The FGT measures available in *PovcalNet* are the headcount index ($\alpha = 0$), the poverty gap (PG) index ($\alpha = 1$) and the squared poverty gap (SPG) index ($\alpha = 2$).

²⁵ The version of the data set used here is for November 2014.

4. Estimating the consumption floor

With a sound sampling design and large enough samples we can be confident about our estimate of the overall mean consumption from a survey for the relevant period. But it is far less clear how reliably we can estimate the consumption floor—the lower bound of the distribution of permanent consumption. If we knew permanent consumptions we could confidently estimate the floor directly. However, that is not the case with the data available. As noted, there are transient consumption shortfalls, whereby recorded consumption for the recall period of the survey interview is temporarily below the floor, but recovers later. Those observed to be very poor are not quite so poor normally. So we cannot credibly estimate the consumption floor as the lowest observed consumption. How might we estimate the floor?

One possible approach would be to set up an explicit parametric model of the joint distribution of permanent consumption and observed consumption. This would allow an estimate of the lower bound, and many other parameters of the marginal distribution of latent permanent consumption. The results would depend, of course, on the assumptions made about the joint distribution, notably about the sources of transient consumption, including measurement errors. There would be many possible solutions depending on those assumptions.

The approach adopted here does not require a complete model of the joint distribution, but focuses instead on the parameter of interest. Of course, assumptions are still required. But the proposed estimator starts with assumptions about the key factor determining the expected value of the floor, namely the probability of any observed consumption being the floor. This allows a simple, transparent, estimator. However, it is acknowledged that the present approach does not say anything about how the estimator might be derived from assumptions about the properties of the underlying joint distribution of permanent and observed consumption.

It is postulated here that observed current consumption equals the floor for permanent consumption at some point within a range of low observed consumptions. A monotonicity property is shown to hold under seemingly weak assumptions, such that the lower observed consumption the higher the probability that this is the consumption floor. The expected value of

the floor is then estimated as a probability-weighted mean of the observed consumptions of a poor stratum, with highest weight on the poorest.²⁶

To formalize this approach, let y denote the n -vector of observed consumptions, with lowest value y^{\min} and highest y^{\max} , while x denotes the corresponding distribution of permanent consumptions for which the lowest level is x^{\min} which is assumed to lie within (y^{\min}, y^{\max}) . Permanent consumption is unobserved, and x^{\min} is treated as a random variable. The task is to use the data on y to estimate $E(x^{\min}|y)$. We can write:

$$E(x^{\min}|y) = \sum_{i=1}^n \phi(y_i) y_i \quad (1)$$

Here $\phi(y_i) = \Pr(y_i = x^{\min})$ is the probability that the observed current consumption of person i , y_i , is in fact the consumption floor x^{\min} . The weights in (1) are not ethically-motivated distributional weights but reflect the probabilities attached to any observed income being in fact the lowest level of permanent consumption. The probabilities are not data, of course.

What properties might we expect $\phi(y_i)$ to have? It seems reasonable to assume that there is an upper-bound to the possible values for the floor. Specifically, it is assumed that beyond some critical level y^* of observed consumptions there is no chance of that being the floor to permanent consumption; $\phi(y_i) = 0$ if and only if $y_i > y^*$. However, it is never certain that any observed y_i is in fact the lower bound to permanent consumption. Additionally, individual consumptions are treated as being distributed independently across people.

It follows that the probability of an observed level of consumption being the true lower bound of permanent consumption falls as observed consumption rises until y^* is reached; i.e., $\phi(y_i)$ is a strictly decreasing function for $y_i \leq y^*$. This property is easily verified on noting that $\Pr(x_j > y_i) = 1 - F_j(y_i)$ for person $j=1, \dots, n$ where $F_j(y_i)$ is the cumulative distribution function for the permanent consumption of j evaluated at the observed consumption y_i , with density function $f_j(y_i) = F'_j(y_i) > 0$. Given independence, we then have:

²⁶ Notice that the use of a group average here is motivated by the data problem, namely that we do not observe the lowest level of permanent consumption. The aim is to estimate that level with the (imperfect) data on hand.

$$\phi(y_i) = \prod_{j=1}^n (1 - F_j(y_i)) \quad (2)$$

On differentiating (and noting that $0 < F_j(y_i) < 1$) we have:

$$\phi'(y_i) = -\phi(y_i) \sum_{j=1}^n \frac{f_j(y_i)}{(1 - F_j(y_i))} < 0 \text{ for } y_i \leq y^* \quad (3)$$

The extent of inequality amongst those with $y_i \leq y^*$ can also be expected to play a role. The expected value of the floor cannot exceed the (un-weighted) mean of observed consumptions for those living under y^* . Imagine that all those living below y^* have the same observed consumption, the mean \bar{y}^* for the q persons with $y_i \leq y^*$. Then it is reasonable to treat \bar{y}^* as the floor (assuming zero-mean error). Now introduce inequality amongst the poor. This implies a larger spread of y 's below the mean and hence a lower $E(x^{\min}|y)$ relative to \bar{y}^* , given that lower observed y 's are more likely to be near the floor. Inequality amongst the poor is reflected in those poverty measures satisfying the transfer axiom, as satisfied by the squared poverty gap index (Section 3). Let P_2^* denote the value of this index when $z = y^*$. Intuitively, we expect a higher P_2^* to be associated with a lower expected floor for any given \bar{y}^* .

To derive an operational measure, the probabilities are assumed to take the form:

$$\begin{aligned} \phi(y_i) &= k(1 - y_i / y^*)^\alpha \text{ for } \alpha \geq 1 \text{ and } y_i \leq y^* \\ &= 0 \text{ for } y_i > y^* \end{aligned} \quad (4)$$

For expositional purposes consider first the special case $\alpha = 1$. To assure that the probabilities sum to unity we also require that $k = 1/(nP_1^*)$ where P_1^* is the PG index for a poverty line of y^* . Thus $\phi(y_i)$ is person i 's share of the aggregate poverty gap treating y^* as the poverty line. The expected value of the floor relative to y^* is a weighted mean of the values of the y_i / y^* (for $y_i \leq y^*$) with weights given by each person's share of the aggregate gap:

$$E(x^{\min}|y) / y^* = \sum_{y_i \leq y^*} \phi(y_i) y_i / y^* \quad (5)$$

Next, consider the value of P_2^* / P_1^* . By construction, this is a weighted mean of the values of $1 - y_i / y^*$ conditional on $y_i \leq y^*$, also with weights given by the shares of the poverty gap:

$$P_2^* / P_1^* = \sum_{y_i \leq y^*} \phi(y_i) (1 - y_i / y^*) \quad (6)$$

Comparing (5) and (6) we immediately have the following formula for the expected value of the consumption floor (in \$'s per person per day):

$$E_1(x^{\min} | y) = y^* (1 - P_2^* / P_1^*) \quad (7)$$

(The subscript "1" is used to indicate that this is the formula for $\alpha = 1$.)

It is plain from (7) that the necessary and sufficient condition for a rising floor is that the proportionate rate of decline in PG exceeds that for SPG when using y^* as the poverty line. Intuitively, a rising floor requires faster progress against the distribution-sensitive SPG measure when based on the observed consumptions. If both poverty measures are falling then one requires that SPG is falling faster than PG for the expected value of the floor to rise.

While the formula in (7) makes the relationship between the expected floor and poverty measures clear, it is still not obvious what role is played by inequality amongst the poor. Linear probabilities allow a straightforward decomposition. Note that $P_1^* = P_0^* (1 - y_i / y^*)$ and that:

$$P_2^* = \sum_{y_i \leq y^*} (1 - y_i / y^*)^2 / n = P_0^* \left[(1 - y_i / y^*)^2 + \sigma^{*2} / y^{*2} \right] \quad (8)$$

where $\sigma^{*2} = \sum_{y_i \leq y^*} (y_i - \bar{y}^*)^2 / q$ is the sample variance amongst those for whom $y_i \leq y^*$. Equation

(7) can then be written as:

$$E(x^{\min} | y) = \bar{y}^* - \frac{\sigma^{*2}}{y^* - \bar{y}^*} \quad (9)$$

This makes clear how the gap between \bar{y}^* and $E(x^{\min} | y)$ reflects the inequality amongst those with $y_i \leq y^*$, as measured by their variance of consumption normalized by the mean gap, $y^* - \bar{y}^*$.

More generally, (4) implies that:

$$E_\alpha(x^{\min} | y) = y^* (1 - P_{\alpha+1}^* / P_\alpha^*) \quad (10)$$

Where P_α^* is the FGT class of measures for $z = y^*$. Note that, while the FGT measures naturally emerge analytically, the interpretation of the parameter α is very different. Here α determines how the probability of being the poorest person falls as observed consumption increases, rather than the degree of aversion to inequality amongst the poor, as in the FGT index.

The parameter value $\alpha = 0$ can be ruled out, as it implies that observed consumptions contain no information about permanent consumption; using $\alpha = 0$ every consumption below y^* is equally likely to be the lowest, so $y^*(1 - P_1/P_0)$ is the mean consumption of those with $y_i \leq y^*$. Linear probabilities ($\alpha = 1$) avoid this problem, and are computationally convenient given that it can be estimated from available data sources, recalling that *PovcalNet* only gives P_α for $\alpha = 0, 1, 2$. However, values of $\alpha > 1$ can be defended, allowing the probability to decline non-linearly as a convex function.²⁷ Testing sensitivity to relaxing linearity in favor of $\alpha = 2$ will require an approximation given that P_3 is not available in *PovcalNet*. A good approximation can be obtained by first noting that:

$$P_3 = P_2 + P_1 - P_0 + \sum_{y \leq y^*} h(y_i)/n \quad (11)$$

where $h(y_i) \equiv 2(y_i/y^*)^2 - (y_i/y^*)^3$. On expanding $h(y_i)$ using a second-degree Taylor series and evaluating at $y_i = \bar{y}^*$ (the mean for $y_i \leq y^*$), the approximation formula used below is:²⁸

$$P_3 \cong P_2 + P_1 + P_0[h(\bar{y}^*) - 1 + h''(\bar{y}^*)\sigma^{*2}/2] \quad (12)$$

(Note that the term in the Taylor series expansion for $h'(\bar{y}^*)$ drops out on aggregation.)

One concern in implementing this approach is selective mortality and fertility, whereby those living closer to the floor are less likely to survive but may well have higher fertility rates.²⁹ The net effect on estimates of the floor is unclear. The following calculations ignore the problem but this is something that might be explored further in future research.

²⁷ For example, for a uniform density $f_j(y_i)$ in (3), $\phi(y_i)$ is automatically a decreasing convex function.

²⁸ I tested the approximation using micro data for Bihar India in 2010. The actual P_3 was 0.022; the approximation using (12) gave 0.023.

²⁹ Ravallion (2016, Chapter 7) reviews the evidence on the economic gradients in demographic variables.

5. Estimates of the expected value of the consumption floor

In implementing the approach outlined above, one can take either an absolute or relative approach to setting y^* . The former approach sets y^* at a constant value in real terms, while the latter fixes instead the proportion of the population who could be living at the floor. However, it does not seem plausible that the same proportion of the population could be living at the floor in a poor society as a rich one; it is more believable that the poorer the society the larger the set of people who could be living at the floor if we knew their true permanent consumptions.

Using the absolute approach, a plausible assumption is to set for y^* according to the national poverty lines found in the poorest countries. This is one of the methods used by Ravallion et al. (2009) to set the international poverty line of \$1.25 a day. So the first key assumption made here is that there is no chance that any observed consumption level above \$1.25 a day corresponds to a true level of consumption that is in fact the floor. The \$1.25 line corresponds to the 20th percentile in 2010. So this is quite a wide range. I test sensitivity to using a lower value for y^* of \$1.00 a day and using a relative definition, such that a constant percentage of the population is identified as the group of people who may be living at the true floor.

Table 1 gives my estimates of the expected value of the floor from the data described in Section 2. Columns 1-3 give the estimates for $y^* = \$1.00$ as well as \$1.25 under linearity of the probabilities ($\alpha = 1$). Column 2 gives the estimated floor calculated using only consumption surveys (dropping all income surveys). Column 4 gives the estimates for $y^* = \$1.25$ with $\alpha = 2$, allowing the probability of an observed consumption being the true floor to be a strictly convex decreasing function of the level of observed consumption. Since values of P_3^* are not data, the calculation of equation (10) for $\alpha = 2$ uses the approximation in (12).

The mean estimated floor is \$0.67 per day (\$0.68 using consumption surveys only) for linear probabilities.³⁰ Using a quadratic function instead the mean falls to \$0.60 per day. Figure 3 plots the estimated consumption floor for $y^* = \$1.25$ using the full sample over 1981-2011, as well as the mean consumptions of both the poor and the overall population of the developing

³⁰ This is the un-weighted mean over time. The inter-temporal variance is so low that it is unlikely that population weighting would make any detectable difference.

world. Panel (b) gives a “blow-up” of the lower portion, also identifying the contribution of inequality amongst the poor, i.e., $\sigma^{*2}/(y^* - \bar{y}^*)$ (recalling equation (9)).

The main source of statistical imprecision in this estimate of the consumption floor is the cut-off point y^* . The global sample sizes for estimating the poverty measures are huge (over 2 million sampled households from over 900 surveys for the recent years, though less as one goes back in time). Using the Ravallion et al. (2009) estimate of the standard error of the \$1.25 a day poverty line, the implied standard error of the present estimate of the floor for $\alpha = 1$ is \$0.10 per day.³¹ The 95% confidence interval for the consumption floor is thus \$0.47 to \$0.87 per day.

It should be recalled that this assumes that there is zero probability of an observed consumption above \$1.25 a day corresponding to the floor. Naturally, a higher (lower) y^* will raise (lower) the estimated floor. If anything, I suspect that \$1.25 is on the high side. Alternatively, if one sets $y^* = \$1.00$ then the time mean of the floor falls to \$0.55.³²

Also notice that this estimation method does not of course require that nobody should be found living below the expected consumption floor based on their current consumption. That would be too stringent given that there will invariably be some people temporarily living below any consumption floor. For 2011, *PovcalNet* indicates that 3.7% of the population of the developing world lived below \$0.67 a day. The proportion living below the lower bound of the 95% confidence interval for this estimate of the floor is 1.8%.³³

It is evident from Figure 3 that the floor has been quite stable over time. The estimated consumption floor rose by only 9 cents per day over 30 years, from \$0.59 to \$0.68, reflecting a (slightly) steeper pace of decline in SPG^* and PG^* . The contribution of inequality amongst those living below \$1.25 rose from \$0.14 to \$0.20 over the period (Table 1, Column 5), representing 19% and 23% of \bar{y}^* respectively.

³¹ Ravallion et al. (2009) used Hansen’s (2002) estimator for a piece-wise linear (“threshold”) model in estimating the relationship between national poverty lines and private consumption per person.

³² One might assume instead that nobody above median could be living at the floor. Then $y^* = \$2.00$ per day (for 2005), giving a time-mean of the floor of \$0.90. However, the median seems an implausibly high value for y^* .

³³ This is probably an overestimate given that *PovcalNet* uses grouped data for many countries, which require curve fitting; the software uses parameterized Lorenz curves fitted to the grouped data. These will give non-zero estimates to very low levels, even when the micro data do not indicate any observations.

The growth rate in the floor (regression coefficient of $\ln \hat{E}_1(x^{\min}|y)$ on time) is 0.34% per annum, with a standard error of 0.08%; using consumption surveys alone, the trend coefficient is 0.21% per annum (s.e.=0.03%). If instead one uses the estimated quadratic function for the probabilities (regressing $\ln \hat{E}_2(x^{\min}|y)$ on time) the trend is 0.48% per year (s.e.=0.14%).

There is a marked divergence between the mean for the poor as a whole and the estimated floors, with a growth rate for the former of 0.46% per annum (s.e.=0.06). (And the divergence is statistically significant; t-test=4.39; prob.=0.14%, using $\alpha = 1$.) Using an upper bound of \$1.00 a day there is even less sign of a positive trend in the implied floor; the estimate of the floor rises from \$0.52 to \$0.53, although it rises then falls (Table 1).³⁴

One possible response to the apparent lack of progress in raising the consumption floor over the last 30 years might be to point to gains in other determinants of human welfare (such as improved health), such that some composite welfare index has shown a gain for the poorest in terms of that index. That may well be so, but those “non-income” gains have no doubt been enjoyed no less by others living above the consumption floor or the mean for the poor. And for them the consumption gains have been far greater. Indeed, the divergence between the mean for the poor and the expected consumption floor is minor compared to the expanding gap between both and the overall mean of household consumption per person, which grew at an annual (per capita) rate of 2.1% over this period (s.e.=0.24%) and the rate of growth roughly doubled from the turn of the century. There is no sign that the upsurge in average living standards in the developing world since 2000 came with upward pressure on the floor (Figure 3). In relative terms, the consumption floor has fallen from 22% of the overall mean in 1981 to 13% in 2011.

Within these “global” aggregates there is considerable variation across countries in both the estimated levels of the floor and their changes over time. Focusing on the countries with at least two surveys in the data set, Figure 4 plots the estimated consumption floors at country level for the latest date against the earliest.³⁵ (The time period varies between observations, from four years to 33 years.) We see that some countries starting out with low floors saw sizable gains over time, although we cannot know how much of this is due to measurement error (an initial underestimation of the floor is followed by an upward correction).

³⁴ The trend coefficient is very close to zero (a coefficient of 0.0002, with a standard error of 0.0009).

³⁵ Observations confined to cases with a positive count of those living below \$1.25 a day.

The two most populous countries have seen the expected value of the floor rising at an above-average rate. In 1981, the floor was \$0.55 a day in China, and rose to \$0.80 in 2008, with a trend (again based on the regression of the log floor on time) of 0.95% per annum (s.e.=0.12; n=28). India has seen the floor rise from around \$0.72 a day in 1982 to \$0.87 in 2012. For India, the National Sample Surveys allow one to construct an unusually long time series back to the early 1950s. The floor was \$0.62 in 1955 and averaged \$0.64 for the 1950s.³⁶ The trend rate of increase over 1955-2012 was 0.54% per annum (s.e.=0.02, n=47).

It is of interest to compare the experience of the developing world over the last 30 years with that of today's rich world since the mid-19th century. Naturally the historical data are sparse and often of questionable quality, so any quantitative assessment must be considered broadly indicative at best. Bourguignon and Morrisson (2002) have compiled distributional data over the period 1820-1992 and merged with Maddison's (1995) estimates of GDP per capita, also back to 1820. Bourguignon and Morrisson only calculated poverty measures for the world as a whole. However, using their data base (which they kindly provided) I calculated poverty measures for those countries considered rich countries today, using the Bourguignon-Morrisson "extreme poverty" line.³⁷ I find that by 1992 the poverty rate had fallen to zero in most of today's rich countries.³⁸ Indeed, it had probably reached zero in most countries by 1960. Yet in the mid-19th century these countries had an average poverty rate of 55% and the level of the consumption floor was 48% of the poverty line on average.³⁹

So the floor more than doubled in today's rich world over the time it escaped extreme poverty. The pace of progress in raising the floor in today's developing is appreciably lower than this. Conservatively, let us assume that the level of the floor doubled in today's rich world over 100 years. The annualized rate of growth is then 0.7%, about double the rate for the developing world over the last 30 years, and also higher than India's rate, though less than China's.

³⁶ There is extra noise in the first few survey rounds when a number of survey design and implementation issues were still being resolved. So some averaging is needed.

³⁷ That line was chosen to synchronize with the poverty rate for 1990 implied by the Chen and Ravallion (2001) "\$1 a day" line. For further discussion see Ravallion (2014a).

³⁸ This is true of Australia, Austria, Belgium, Canada, Czechoslovakia, Germany, Japan, Hungary, Korea, Luxembourg, New Zealand, the Scandinavian countries, Switzerland and Taiwan. Ravallion (2014a) gives the results. Exceptions are the United Kingdom and the United States, with 2.5% and 3.0% still living in extreme poverty by 1992. While not strictly comparable, the estimates in Ravallion and Chen (2013b) suggest that at some time after 1990 the number of people living below \$1.25 a day in the high-income countries went to zero.

³⁹ In the country groupings used by Bourguignon and Morrisson the consumption floor in 1850 was 45% of the poverty line in Australia-Canada-New Zealand, 51% in Austria-Czechoslovakia-Hungary, 54% in Belgium-Luxembourg-Switzerland, 55% in Germany, 41% in Japan, 42% in Korea-Taiwan and 49% in Scandinavia.

The above results have used a fixed absolute standard for setting y^* . It was argued that this is more plausible than a relative approach to defining the stratum of people who could be living at the floor. However, it should be noted that using a relative standard implies a rising absolute floor over time. For example, suppose one focuses instead on the poorest 20%, corresponding closely to the absolute standard of \$1.25 a day in 2010. If one defines the group of people who are potentially living at the floor in 1981 as the poorest 20% then the estimate of $E_1(x^{\min}|y)$ falls to \$0.37 a day, with a value of y^* for that year of \$0.63 (only slightly higher than the estimate of $E_1(x^{\min}|y)$ using $y^* = \$1.25$). This suggests far greater progress in raising the floor than the absolute approach, with its value almost doubling over 30 years. The “relative floor” has remained a fairly constant % of the overall mean (14% in 1981 and 13% in 2011). The bulk of the drop in the floor for 1981 using the relative definition for y^* is due to the fact that the relative bound has almost halved; the value of SPG/PG is not much different between the two approaches (0.53 in 1981 using the absolute approach versus 0.42 using the relative approach).

6. Validation using other evidence

As Section 4 emphasized, the fact that permanent consumption is unobserved means that the consumption floor must be estimated from the observed data under certain assumptions. Then it is of interest to see whether these results can be confirmed from other sources and using different assumptions. The rest of this paper points to three pieces of evidence.

One source is the largely independent estimate that Lindgren (2015) has made of the biological floor. Recall that the consumption floor is not defined here as the biological minimum, but as the lower bound for permanent consumption. So the present paper’s estimate of the consumption floor should not be lower than Lindgren’s “physical minimum line,” which is the cost of a “barebones basket” of food items assuring 2100 calories per person per day plus a minimal allowance for essential nonfood items. To the nearest cent, Lindgren’s estimate of the biological floor is the same as the present paper’s estimate of the mean floor of \$0.67 per day under linearity, and only slightly higher than for the quadratic function. This suggests that the developing world has not had much success in raising the consumption floor above the biological floor.

A second source of evidence is found in national poverty lines. These can be thought of as the sum of two components: an absolute consumption floor plus a relative component that depends positively on the country's mean consumption. This suggests an alternative method of defining the floor as the expected value of the national poverty line at zero mean. Two sources of data are available to implement this alternative method. Ravallion et al. (2009) compiled national poverty lines for 75 countries, but this is limited to one observation per country. A second source was developed by Jolliffe and Prydz (2015). These were derived from national poverty measures in World Bank *Poverty Assessments*, by solving for the line that yields the national poverty rate, taken as given. The Jolliffe-Prydz data provide a much larger sample of over 600 lines, with the advantage of multiple observations over time for many countries. Using their estimates (kindly provided by Jolliffe and Prydz) I find:⁴⁰

$$z_i = \underset{(0.072)}{0.661} + \underset{(0.011)}{0.395} \bar{y}_i + \hat{\varepsilon}_i \quad R^2=0.872, n=609 \quad (13)$$

Figure 5 plots the data and the regression line. The implied floor of \$0.66 per day is very close to the estimates from Section 5. There is very little obvious sign of nonlinearity in the relationship in Figure 5, and this was confirmed by adding higher-order polynomials (up to order 4) in the mean to (13);⁴¹ the estimates of the floor varied, but none were significantly different to the estimates from Section 5.

As we saw in Figure 3, the paper's main estimate of the floor shows much less growth over time than found in the overall mean. Maybe there is some source of bias that is leading my method to underestimate how much the floor is growing with overall growth. It is thus of interest to see if this finding is also robust to the use of this alternative data source. The national lines were set at different dates. On adding a time trend to the regression in (13) one finds no significant drift in the consumption floor; the regression coefficient on time is 0.010 (s.e.=0.007).

The third source of supporting evidence is found in the traditional counting approach, to which we now turn.

⁴⁰ White standard errors in parentheses. Using the smaller data set from Ravallion et al. (2009) the corresponding estimate of the floor was \$0.647 (s.e.=0.288; n=73).

⁴¹ The preferred functional form in Ravallion et al. (2009) is not feasible since it uses a log transformation of \bar{y}_i for which $E(z_i | \bar{y}_i = 0)$ is undefined.

7. Revisiting the counting approach

The traditional counting approach suggests substantial gains to the poor of the developing world over the last 30 years. Figure 6 gives the cumulative distribution functions (CDF's) for 1981 and 2011 in the upper panel, and the vertical differences between the CDFs in the lower panel.⁴² We see that there is first-order dominance, implying an unambiguous reduction in poverty for all possible lines and all additive measures, as was found by Chen and Ravallion (2010) for a shorter period.⁴³

However, there is a feature of Figure 6 that has not (to my knowledge) been noted before, but is of interest from the point of view of corroborating the evidence above that the floor has risen little over these 30 years. Figure 7 makes the point clearer by giving the monetary gain at each percentile implied by Figure 6, i.e., the absolute difference between the quantile functions, obtained by inverting the CDFs.⁴⁴ (These gains are simply the horizontal differences between the CDFs in Figure 6.) Consistently with the lack of progress in raising the floor we see that the gains are close to zero for the poorest, but rising to quite high levels. This is also consistent with what we know about rising absolute inequality in the developing world (Ravallion, 2014c).

A further insight from Figure 7 is that there are larger absolute gains for the second decile from the bottom (though fairly flat between the 10th and 20th percentiles). Using the 20th percentile as the cut-off point in the relative approach is thus picking up these gains. At a sufficiently low cut-off, even the relative approach will show little gain in the floor.

Might the counting approach pick up the lack of progress for the poorest if one looks well below the \$1.25 line? Figure 8 gives the poverty rates for the developing world for various lines, all of which indicate a reasonably steady decline over time. To provide a simple measure of the incidence of “ultra-poverty,” let us focus on the \$0.87 a day, which is the upper bound of the 95% confidence interval for the estimated consumption floor, as described in Section 4.⁴⁵ We see in Figure 8 that this has declined steadily over time in the developing world as a whole. This also holds for most countries, as can be seen in Figure 9, comparing the earliest and latest surveys for

⁴² The CDF is truncated above \$20 a day to give greater detail at the lower end; however, there is first-order dominance over the entire range.

⁴³ On the implications of first-order dominance in this context see Atkinson (1987).

⁴⁴ The empirical quantile function is used for 1981. For the purpose of creating the graph, the quantile function for 2011 was based on a 10th degree polynomial, which fitted extremely well ($R^2=0.998$), although the top 2% were trimmed as these are considered less reliable.

⁴⁵ The use of the 95% confidence interval is essentially arbitrary. I also give results for other lines.

those countries in *PovcalNet* with two or more surveys. The number of people living in ultra-poverty by this definition fell from 1,317 (35.4%) million to 423 million (7.1%) over this period. (Using the more stringent definition of \$0.77 a day, the percentage declined from 1,098 million (29.6%) to 308 million (5.2%) in 2011.) For the developing world as a whole, the share of total poverty represented by the ultra-poor fell from 67% in 1981 to 42% in 2011. The bulk of the reduction in overall poverty rates (for \$1.25 a day or \$2.00) is accountable to a lower incidence of ultra-poverty. Between 1981 and 2011 the \$1.25 a day poverty rate fell by 35.8% points; almost 80% of this decline (28.4% points) is accountable to the decline in the ultra-poverty rate.

The trend (regression coefficient on time) over 1981-2011 for the percentage of ultra-poor is -0.83% points per annum (with a standard error of 0.07%).⁴⁶ This is lower than the trend for the percentage below \$1.25 a day of -1.13% (s.e.=0.04%), but the difference is not large, with the implication that the bulk of the inter-temporal variance in the overall poverty rate is accountable to progress against ultra-poverty; the R^2 for the regression of the overall poverty rate for \$1.25 and the ultra-poverty rate for \$0.87 is 0.97. Even more strikingly, progress against ultra-poverty also accounts for the bulk of the progress against poverty judged by the \$2.00 line. The poverty rate for the latter line has an annual trend of -1.12% (s.e.=0.09%), almost identical to that for the \$1.25 line. For the \$2.00 line, the R^2 for the regression of the overall poverty rate on the ultra-poverty rate for \$0.87 is 0.91.

This pattern is also evident at country level. Over three-quarters (77.4%) of the variance in annualized rates of poverty reduction using the \$1.25 line is accountable to rates of progress against ultra-poverty. Only 13.6% is accountable to changes in the density of those who were poor but not ultra-poor; the covariance term accounts for 9.0%. Figure 10 plots the rate of change in P_0 for \$1.25 a day across countries against the corresponding change in the ultra-poverty rate. There is close to a 1-to-1 relationship; as the number of ultra-poor in a country falls, we also see roughly similar exit rates from the ranks of the poor population as a whole.

This pattern is suggestive of a process of what can be called rank-preserving lifting out of poverty. It is as though, as one of the group of “poor but not ultra-poor” is lifted out of poverty this frees up space for one of the ultra-poor, who moves up to take that spot on the ladder. But the floor rose very little.

⁴⁶ For the \$0.77 line the annual trend is -0.69% (s.e.=0.05).

8. Why so little progress in raising the floor?

We saw in section 5 that today's high-income countries appear to have been more successful in raising the floor above the biological minimum over the 100 years after 1850 than we have seen in today's developing world over the last 30 years. It is a plausible hypothesis that social protection policies in today's rich world since the early 20th century played a role; the period after World War 1 saw a substantial expansion in such policies in Western Europe, the United Kingdom, North America and Australia, and this continued for the rest of that century.⁴⁷

As noted in the introduction there has also been an expansion in SSN coverage in the developing world. However, the best available evidence suggests that only about one third of those families in the poorest quintile in the developing world are receiving any direct help from existing safety net policies. And the performance tends to be worse in poorer countries. These observations are based on data compiled by the World Bank on the coverage of safety-net programs across the developing world, using household surveys that identified direct beneficiaries of these for each of over 100 countries spanning 1998-2012. Comparing regional averages one finds that the coverage of the poorest quintile is weaker in the two poorest regions, Sub-Saharan Africa and South Asia. In SSA, only 20% of the poorest 20% of the population (ranked by income or consumption per person) receive anything from the social safety net. By contrast, in Latin America the proportion is 53%.⁴⁸ Figure 11 gives the data at country level.

Taking a simple average across countries, the data indicate that only about half (48%) of the poorest quintile receive anything from the public social safety net; on weighting by population the share falls to 36%. However, there is huge variation, spanning the range from virtually zero to virtually 100% coverage. Some of this is undoubtedly measurement error. But there is clearly a strong and positive income gradient across countries in safety-net coverage. The average elasticity of social safety net coverage of the poor to GDP is about 0.9.⁴⁹

It is notable that the coverage rate for the poor tends to exceed that for the population as a whole. The average difference between the two coverage rates is not large, although it tends to

⁴⁷ On the history of social protection policies in the U.K. and U.S. see Mencher (1967).

⁴⁸ See World Bank (2014). For South Asia the overall coverage rate is 25%, for MENA it is 28%, for East Asia it is 48% while for EECA it is 50%.

⁴⁹ The regression coefficient of the log of coverage rate for the poor on the log of GDP per capita is 0.91 with a standard error is 0.13. The corresponding elasticity for the population as a whole is 0.80 (s.e.=0.11). If one controls for the overall coverage rate of the population there is no longer any statistically significant effect of GDP on the coverage rate of the poorest quintile.

rise with GDP per capita.⁵⁰ Richer countries tend to be markedly better at covering their poor, although the bulk of this is explained by differences in the overall coverage rate.

None of this means that poor countries are powerless to help their poor through direct interventions. Indeed, we see in Figure 11 that some low-income countries do quite well. There are also signs that developing countries are doing better in this respect over time. Unfortunately there are only 25 countries with more than one observation in the World Bank database. Comparing the latest and earliest surveys for those countries, I estimate that the overall coverage rate (for the population as a whole) is increasing at 3.5% points per year (standard error of 1.1% points). The coverage rate for the poor is not increasing at quite the same pace; for them the rate of increase is 3.0% points per year (standard error of 1.0%).

9. Conclusions

The popularity and influence of the view that the poorest are left behind in growing developing economies begs for an empirical assessment of its validity. The evidence found in the literature does not appear to be consistent with this view. However, there is an important conceptual difference between focusing on counts of poor people (following in the footsteps of Bowley and others) versus focusing on the level of living of the poorest, in the spirit of Gandhi's talisman or the Rawlsian difference principle. Both perspectives are evident in past thinking and policy discussions. Both have been advocated as development goals, although the counting approach, as implemented in various poverty measures, has dominated the attention of economists and statisticians monitoring progress against poverty. Only when we recognize this conceptual difference in the approaches taken to measuring poverty can we understand why we hear very different answers to the question posed in the title of this paper.

The paper has demonstrated that success in assuring that “no-one is left behind” can be readily monitored from existing data sources under certain assumptions. The proposed approach recognizes that there are transient consumption effects in the observed survey data. However, the data are assumed to be reliable enough to assure that it is more likely that the person with the lower observed consumption is living at the floor than anyone else. The key assumption is that the probability of being the poorest person falls as a power function. Then the expected value of

⁵⁰ Regressing the log of the ratio of coverage rate for the poor to the overall coverage rate on the log of GDP per capita gives a regression coefficient of 0.16, with a standard error of 0.04.

the floor can be derived as a function of standard Foster-Greer-Thorbecke poverty measures. For example, if the probability falls linearly up to the upper bound then the ratio of the squared poverty gap to the poverty gap relative to that bound—two readily-available poverty measures—emerges as the key (inverse) indicator for assessing progress in raising the floor.

Drawing on the results from household surveys for developing countries spanning 1981-2011, the paper finds considerable progress against poverty using the counting approach. There is first-order dominance over the 30 years, implying an unambiguous reduction in absolute poverty by the counting approach over all lines and all additive measures (including distribution-sensitive measures). Mean consumption per capita in the developing world has been growing at around two percent per annum over this period, and four percent since 2000.

However, there appears to have been very little absolute gain for the poorest. The precise estimates of the mean value of the consumption floor are around half of the \$1.25 international poverty line. Using an absolute approach to identifying the floor, the increase in the level of the floor seen over the last 30 years or so has been small—far less than the growth in mean consumption. The modest rise in the mean consumption of the poor has come with rising inequality (specifically, a rising variance normalized by the mean poverty gap), leaving room for only a small gain in the level of living of the poorest. The bulk of the developing world's progress against poverty has been in reducing the number of people living close to the consumption floor, rather than raising the level of that floor. In this sense, it can be said that the poorest have indeed been left behind.

This is consistent with certain models of dynamic poverty traps, whereby some people are stuck in a low-level equilibrium. It is also consistent with the evidence of weak coverage of poor people by existing social safety nets in developing countries. With better social policies it may be conjectured that we will see more progress in lifting the floor in the future, consistently with the evidently faster progress that today's rich world made in the 100 years or so after 1850 in bringing the consumption floor well above the biological floor.

Stronger indications of a rising floor are found if one adopts a relative approach to defining the upper bound on consumption for those people who could conceivably be living at the floor, and one sets the fixed percentage at a sufficiently high level. For example, focusing on the poorest 20% suggests considerable progress in raising the expected value of the floor. However, the paper has argued that an absolute approach makes more sense on the grounds that

one expects a poorer society to have more people living near the floor, as is found to be the case empirically using the counting approach.

To anticipate one response, it might be argued that progress in lifting the floor is a second-order issue, as long as fewer people live near the floor. That is implicit in the traditional counting methods used to assess progress against poverty. However, proponents of this view must surely take pause when one notes that for a long time, and across countries at very different levels of development, social policies have often claimed that they aim to ensure a minimum level of living above any biological consumption floor required for mere survival. Negative income tax schemes and (formally-equivalent) basic-income guarantees financed by progressive income taxes aim to raise society's consumption floor above the biological minimum. And such efforts are not confined to rich countries; indeed, the two largest anti-poverty programs in the world today (in China and India) aim to raise the floor. In forming their views, casual observers may well focus on the observed level of living of those they deem to be the poorest.

While it would be ill-advised to look solely at the level of the floor, it can be acknowledged that this has normative significance independently of attainments in reducing the numbers of people living near that floor. The thesis of this paper is not that progress against poverty should be judged solely by the level of the consumption floor, but only that the latter should no longer be ignored in practice.

References

- Aghion, Philippe, Eve Caroli and Cecilia Garcia-Penalosa, 1999, "Inequality and Economic Growth: The Perspectives of the New Growth Theories," *Journal of Economic Literature* 37(4): 1615-1660.
- Ahmed, Akhter, Ruth Vargas Hill, Lisa C. Smith, Doris M. Wiesmann, and Tim Frankenberger, 2007, *The World's Most Deprived*. Washington DC: International Food Policy Research Institute.
- Atkinson, Anthony B., 1987, "On the Measurement of Poverty," *Econometrica* 55: 749-764.
- Azariadis, Costas, 1996, "The Economics of Poverty Traps. Part One: Complete Markets," *Journal of Economic Growth* 1: 449-486.
- _____, 2006, "The Theory of Poverty Traps: What Have we Learned?," in Samuel Bowles, Steven Durlauf and Karla Hoff (eds) *Poverty Traps*, Princeton: Princeton University Press.
- Banerjee, Abhijit, and Andrew Newman, 1994, "Poverty, Incentives and Development," *American Economic Review* 84(2): 211-215.
- Barrientos, Armando, 2013, *Social Assistance in Developing Countries*. Cambridge: Cambridge University Press.
- Ben-David, Dan, 1998, Convergence Clubs and Subsistence Economies, *Journal of Development Economics* 55: 153–159.
- Blackorby, Charles, and David Donaldson, 1984, "Social Criteria for Evaluating Population Change," *Journal of Public Economics* 25: 13–33.
- Blaug, Mark, 1962, *Economic Theory in Retrospect*, London: Heinemann Books.
- Bowley, Arthur L., 1915, *The Nature and Purpose of the Measurement of Social Phenomena*. London: P.S. King and Sons.
- Chen, Shaohua, and Martin Ravallion, 2010, "The Developing World is Poorer than we Thought, but no Less Successful in the Fight against Poverty," *Quarterly Journal of Economics* 125(4): 1577-1625.
- Cichon, Michael, and Krzysztof Hagemeyer, 2007, "Changing the development policy paradigm: Investing in a Social Security Floor for All," *International Social Security Review* 60: 169-196.

- Dasgupta, Partha, 1993, *An Inquiry into Well-Being and Destitution*, Oxford: Oxford University Press.
- Dasgupta, Partha, and Debraj Ray, 1986, "Inequality as a Determinant of Malnutrition and Unemployment," *Economic Journal* 96: 1011-34.
- Dollar, David, and Kraay, Aart, 2002, "Growth is Good for the Poor," *Journal of Economic Growth* 7(3): 195-225.
- Dollar, David, Tatjana Kleineberg and Aart Kraay, 2013, "Growth Still is Good for the Poor," Policy Research Working Paper 6568, World Bank.
- Dutta, Puja, Rinku Murgai, Martin Ravallion and Dominique van de Walle, 2014, *Right-to-Work? Assessing India's Employment Guarantee Scheme in Bihar*. World Bank.
- Fields, Gary, 2001, *Distribution and Development*, New York: Russell Sage Foundation.
- Fleurbaey, Marc, and François Maniquet, 2011, *A Theory of Fairness and Social Welfare*, Cambridge: Cambridge University Press.
- Freiman, Christopher, 2012, "Why Poverty Matters Most: Towards A Humanitarian Theory Of Social Justice," *Utilitas* 24(1): 26-40.
- Friedman, Milton, 1957, *A Theory of the Consumption Function*, Princeton NJ: Princeton University Press.
- Foster, James, J. Greer, and Erik Thorbecke, 1984, "A Class of Decomposable Poverty Measures," *Econometrica* 52: 761-765.
- Foster, James, Suman Seth, Michael Lokshin, and Zurab Sajaia, 2013, *A Unified Approach to Measuring Poverty and Inequality*. Washington DC: World Bank.
- Galor, Oded, and Joseph Zeira, 1993, "Income Distribution and Macroeconomics." *Review of Economic Studies* 60(1): 35-52.
- Hammond, Peter, 1976, "Equity, Arrow's Conditions and Rawls' Difference Principle," *Econometrica* 44: 793-804.
- Hansen, Bruce E., 2000, "Sample Splitting and Threshold Estimation," *Econometrica* 68: 575-603.
- Harsanyi, John, 1975, "Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls's Theory," *American Political Science Review* 69(2): 594-606.

- International Labor Organization, 2012, "[R202:Social Protection Floors Recommendation](#)," Recommendation concerning National Floors of Social Protection Adoption, Geneva, 101st ILC session (14 June 2012).
- James, Aaron, 2015, "World's poorest "left behind", says Vatican," [Christian Radio](#) 11 February.
- Jolliffe, Dean, and Espen Prydz, 2015, "Global Poverty Goals and Prices. How Purchasing Power Parity Matters," Policy Research Working Paper 7256, World Bank.
- Kraay, Aart, and Claudio Raddatz, 2007, "Poverty Traps, Aid and Growth," *Journal of Development Economics* 82(2): 315-347.
- Lewis, Arthur, 1954, "Economic Development with Unlimited Supplies of Labor," *Manchester School of Economic and Social Studies* 22: 139-191.
- Lindgren, Mattias, 2015, "The Elusive Quest for the Subsistence Line. How Much Does the Cost Of Survival Vary Between Populations?" Comparative Institutional Analysis Working Paper 2015:1, Lund University, Sweden.
- Lipton, Michael, 1988, "The Poor and the Ultra-Poor: Some Interim Findings," Discussion Paper 25, World Bank, Washington DC.
- Lopez, Humberto, and Luis Servén, 2009, "Too Poor to Grow," Policy Research Working Paper 5012, World Bank.
- Malthus, Thomas Robert, 1806, *An Essay on the Principle of Population, 1890 Edition*, London: Ward, Lock and Co.
- Marshall, Thomas, 1950, *Citizenship and Social Class and Other Essays*. Cambridge: Cambridge University Press.
- Mencher, Samuel, 1967, *Poor Law to Poverty Program: Economic Security Policy in Britain and the United States*, Pittsburgh: University of Pittsburgh Press.
- Mirrlees, James, 1975, "A Pure Theory of Underdeveloped Economies," in: L. G. Reynolds (ed.), *Agriculture in Development Theory*, New Haven: Yale University Press.
- Michelbach, Philip A., John T. Scott, Richard E. Matland, and Brian H. Bornstein, 2003, "Doing Rawls Justice: An Experimental Study of Income Distribution Norms," *American Journal of Political Science* 47(3): 523-539.
- Paine, Thomas, 1797, *Agrarian Justice*, 2004 edition published with *Common Sense* by Penguin.

- Radelet, Steven, 2015, *The Great Surge: The Unprecedented Economic and Political Transformation of Developing Countries Around the World*. Simon & Schuster.
- Ravallion, Martin, 1995, "Growth and Poverty: Evidence for Developing Countries in the 1980s," *Economics Letters* 48, 411-417.
- _____, 2001, "Growth, Inequality and Poverty: Looking Beyond Averages." *World Development* 29(11): 1803-1815.
- _____, 2010, "Do Poorer Countries Have Less Capacity for Redistribution?" *Journal of Globalization and Development* 1(2): 1-29.
- _____, 2012, "Why Don't we See Poverty Convergence?" *American Economic Review* 102(1): 504-523.
- _____, 2014a, "Poverty in the Rich World When It Was Not Nearly So Rich," Blog Post, Center for Global Development, Washington DC.
<http://international.cgdev.org/blog/poverty-rich-world-when-it-was-not-nearly-so-rich>.
- _____, 2014b, "An Emerging New Form of Social Protection in 21st Century China," in *Oxford Companion to the Economics of China*, edited by Shenggen Fan, Ravi Kanbur, Shang-jin Wei, Xiaobo Zhang, Oxford: Oxford University Press.
- _____, 2014c, "Income Inequality in the Developing World," *Science* 344: 851-5.
- _____, 2016, *The Economics of Poverty. History, Measurement and Policy*, New York: Oxford University Press.
- Ravallion, Martin, and Shaohua Chen, 2003, "Measuring Pro-Poor Growth," *Economics Letters* 78(1): 93-99.
- _____, 2011, "Weakly Relative Poverty," *Review of Economics and Statistics* 93(4): 1251-1261.
- _____, 2013a, "More Relatively Poor People in a Less Absolutely Poor World," *Review of Income and Wealth* 59(1): 1-28.
- _____, 2013b, "A Proposal for Truly Global Poverty Measures," *Global Policy*, 4(3): 258-265.
- Ravallion, Martin, Shaohua Chen and Prem Sangraula, 2009, "Dollar a Day Revisited," *World Bank Economic Review* 23(2):163-184.
- Rawls, John, 1971, *A Theory of Justice*, Cambridge MA: Harvard University Press.
- Roemer, John, 1996, *Theories of Distributive Justice*, Cambridge: Harvard University Press.

- _____, 2014, "Economic Development as Opportunity Equalization," *World Bank Economic Review* 28(2): 189-209.
- Sen, Amartya, 1976, "Poverty: An Ordinal Approach to Measurement," *Econometrica* 46: 437-446.
- Stiglitz, Joseph E., 1976, "The Efficiency Wage Hypothesis, Surplus Labour and the Distribution Of Income In L.D.Cs," *Oxford Economic Papers* 28: 185-207.
- United Nations, 2011, *The Millennium Development Goals Report*, New York: United Nations Secretariat.
- World Bank, 2014, *The State of Social Safety Nets 2014*, Washington DC: World Bank.

Figure 1: Same reduction in the poverty count but different implications for the poorest

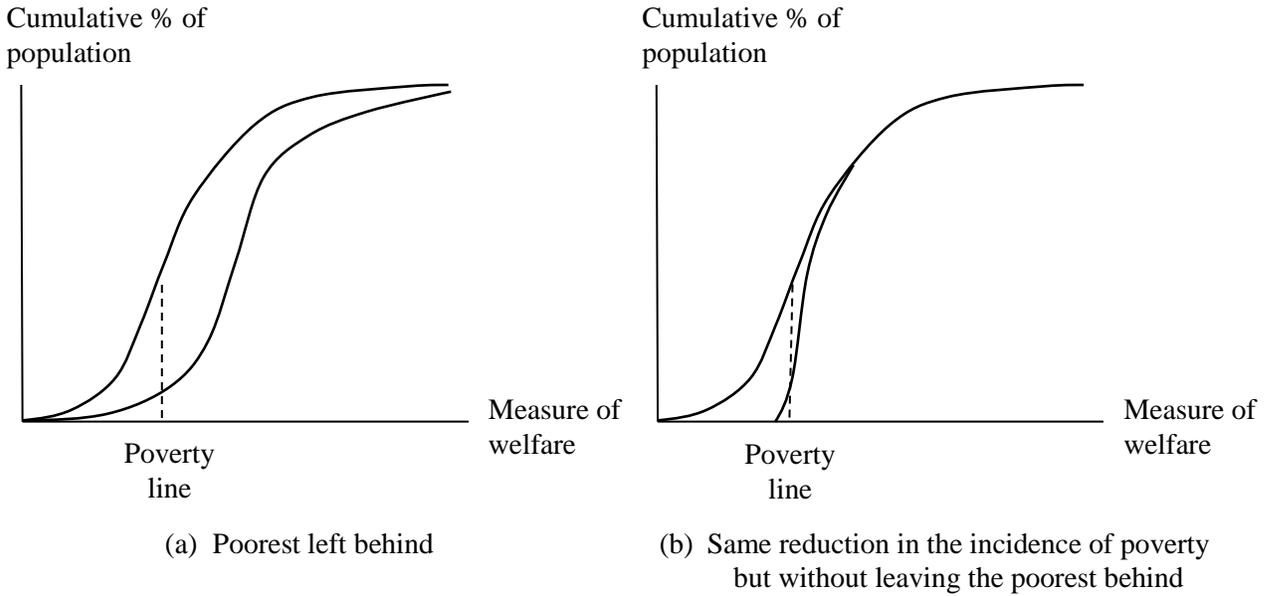


Figure 2: Stylized representation of the relationship between current and permanent consumption

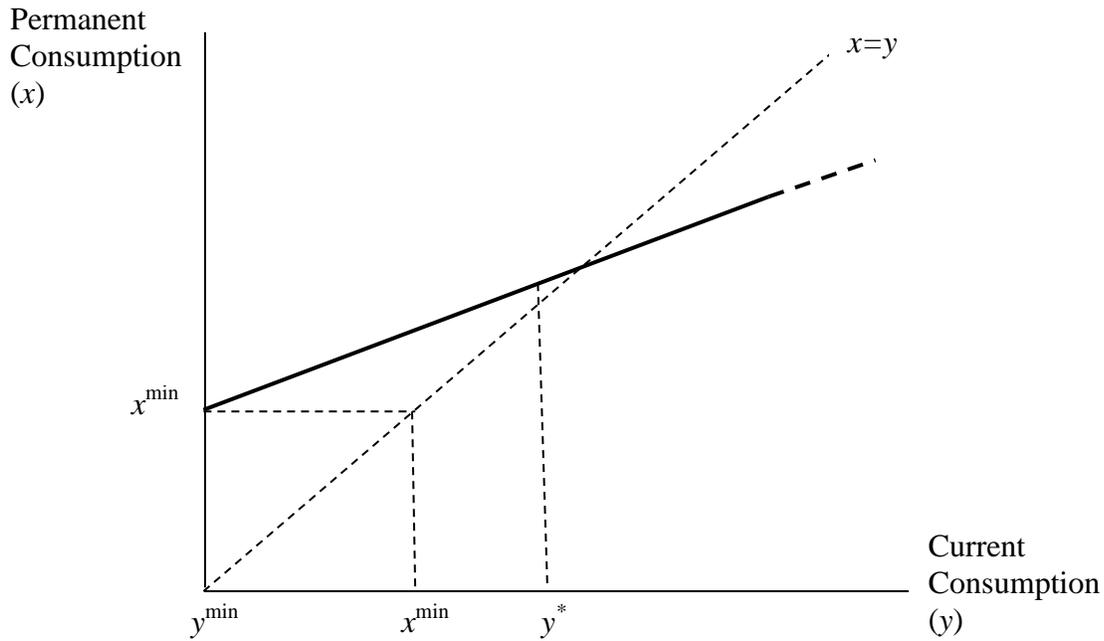
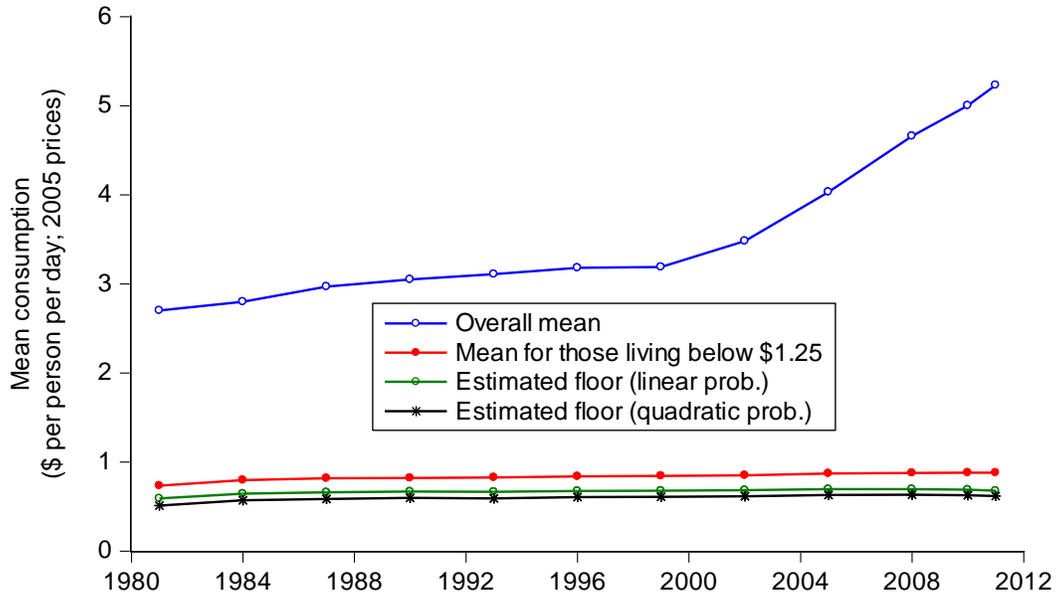


Figure 3: Mean consumptions for the developing world

(a) Including overall mean



(b) Blow up lower segment of panel (a)

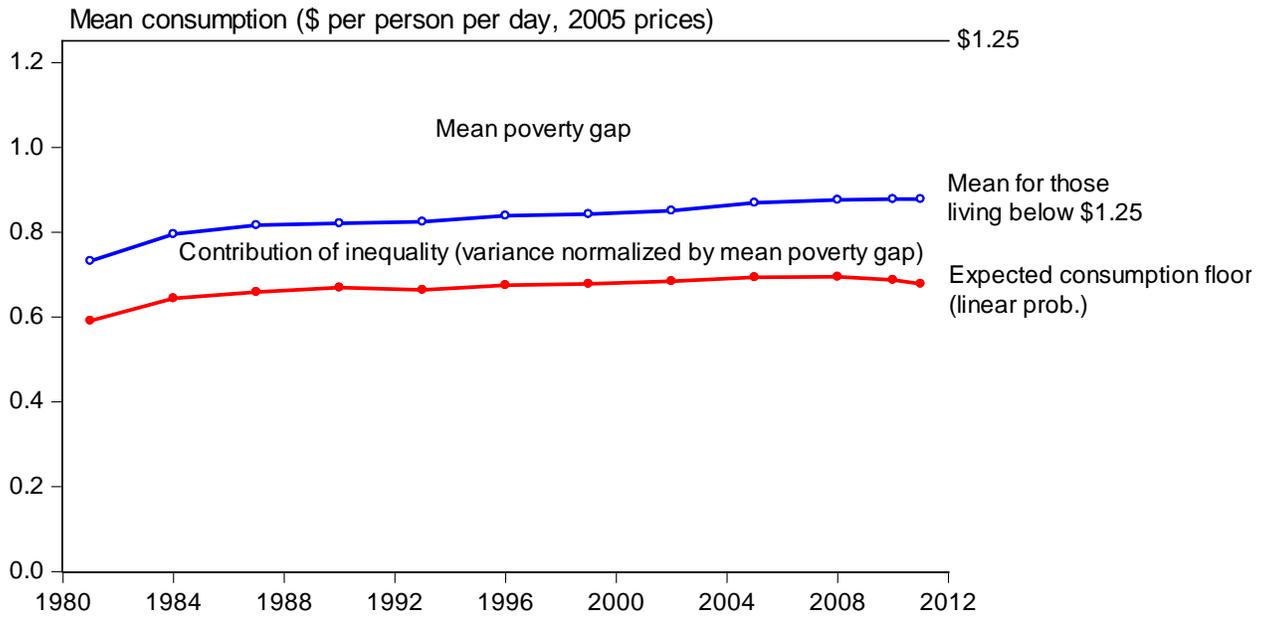


Figure 4: Consumption floors over time across countries

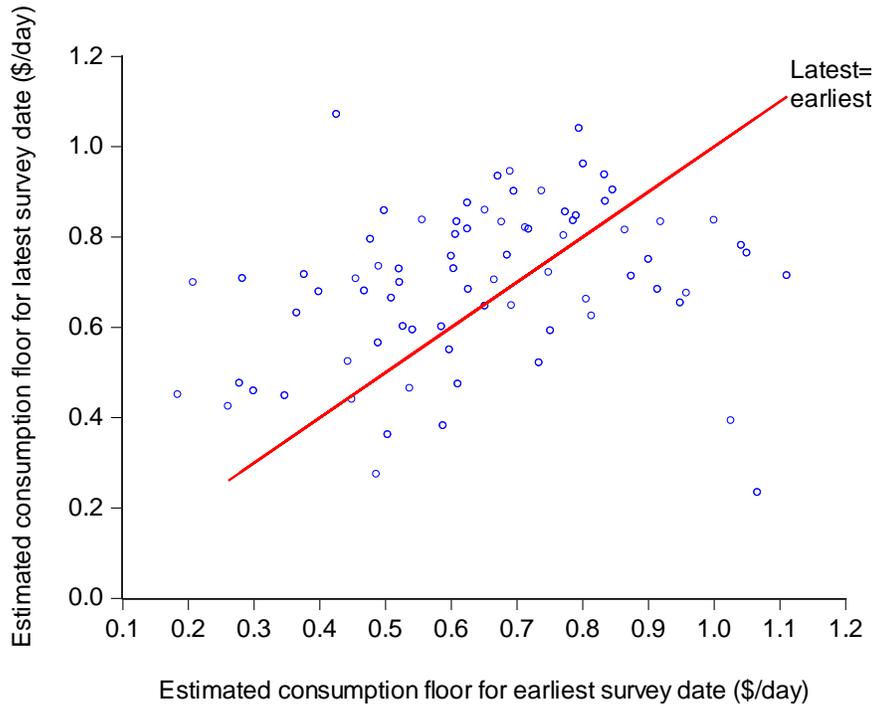


Figure 5: National poverty lines plotted against the closest survey mean

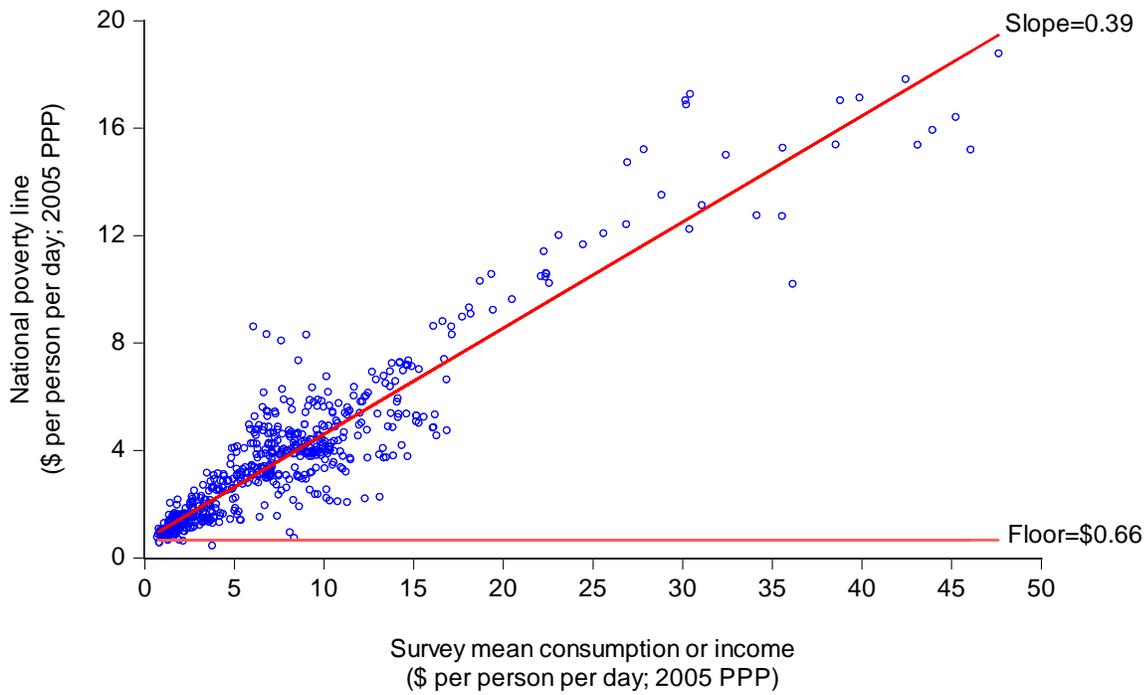


Figure 6: Cumulative distribution functions for the developing world 1981-2011

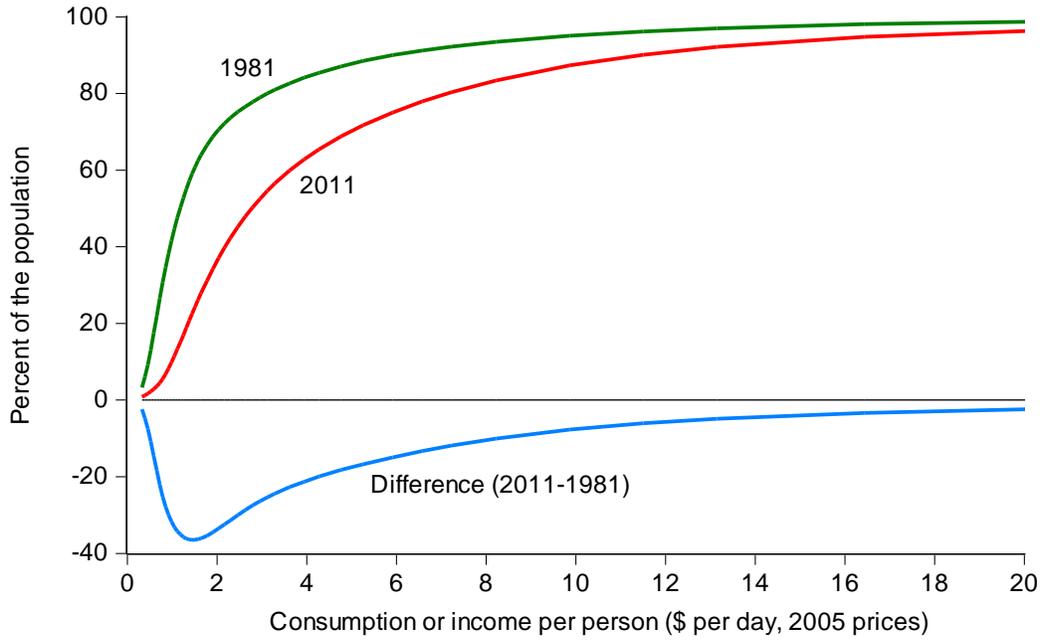


Figure 7: Absolute gains by percentile 1981-2011

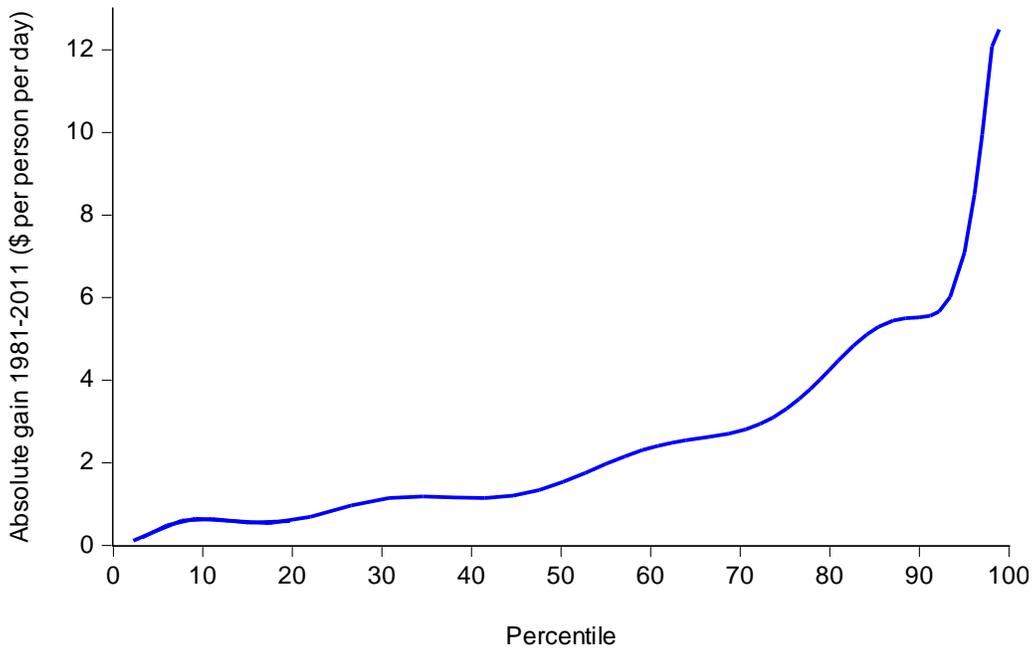


Figure 8: Percentage of the population of the developing world living below each line

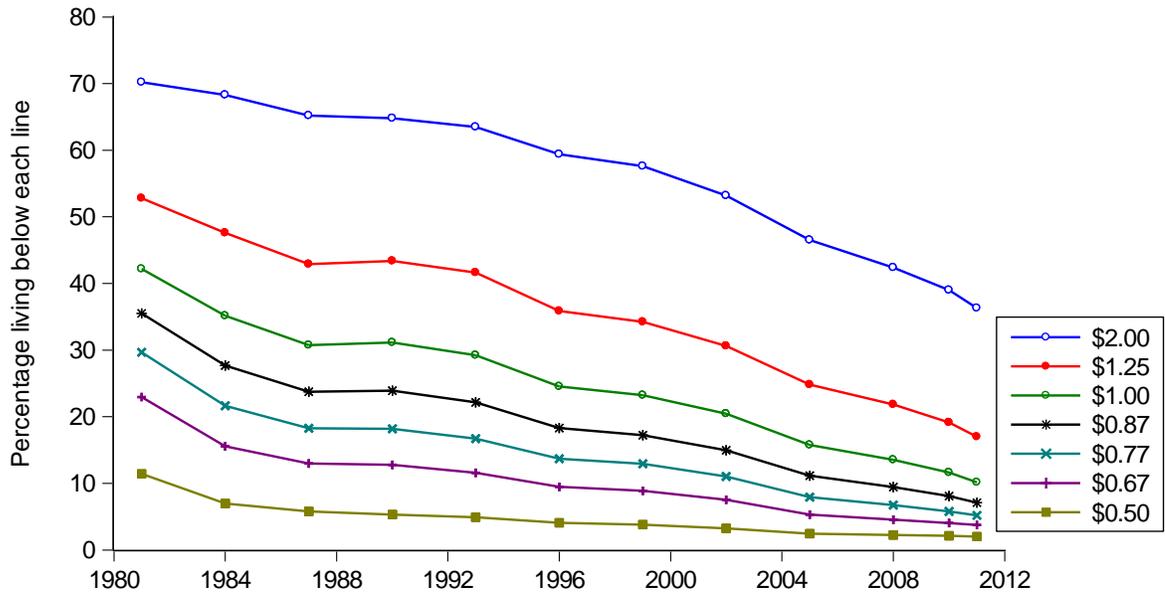


Figure 9: Changes in the incidence of ultra-poverty at country level

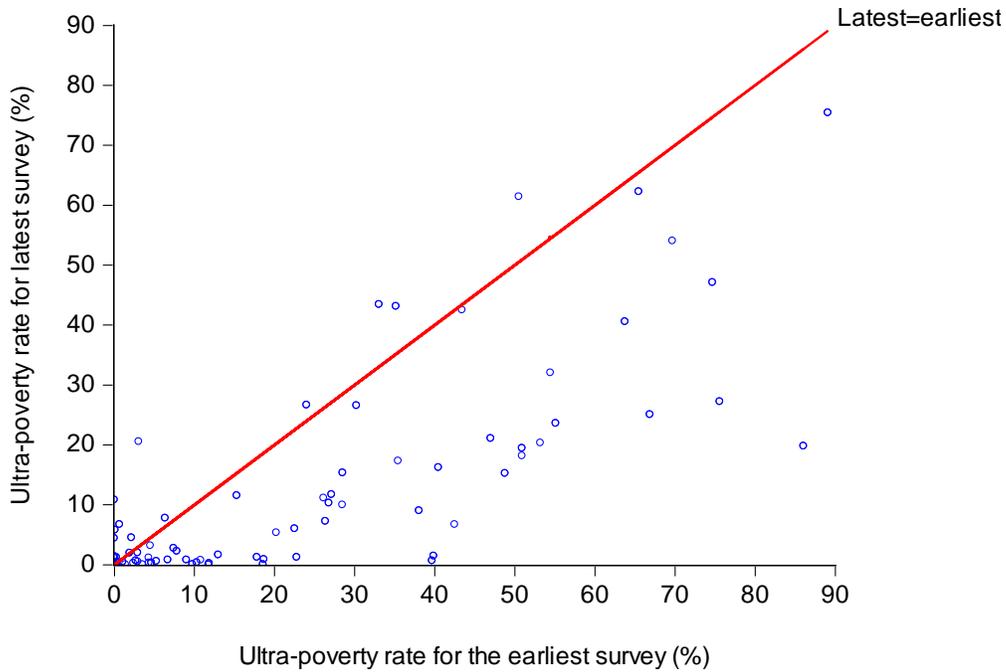


Figure 10: Progress against ultra-poverty at country level translated into progress against total poverty

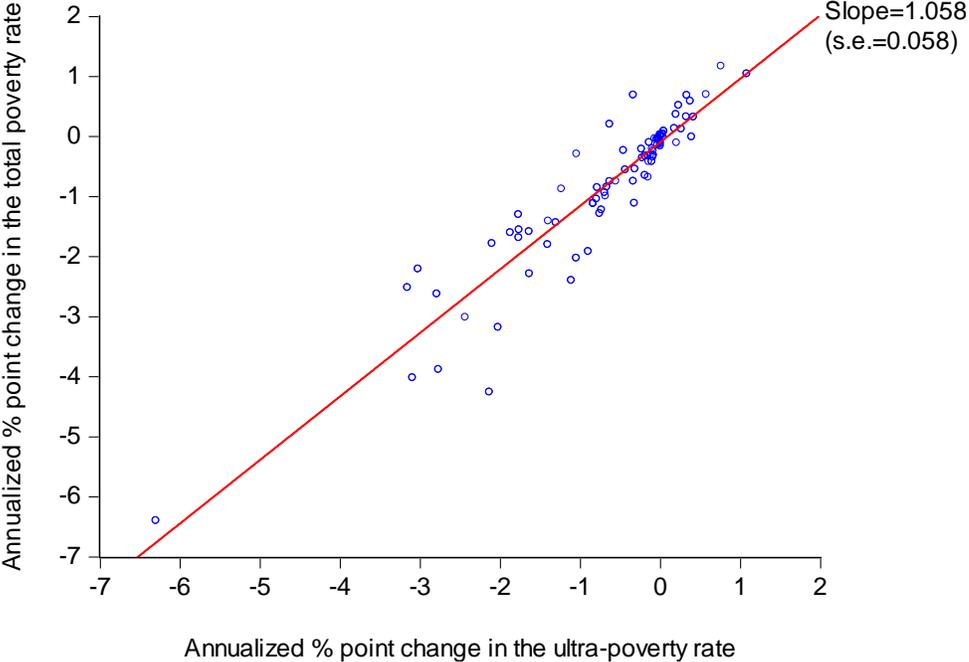
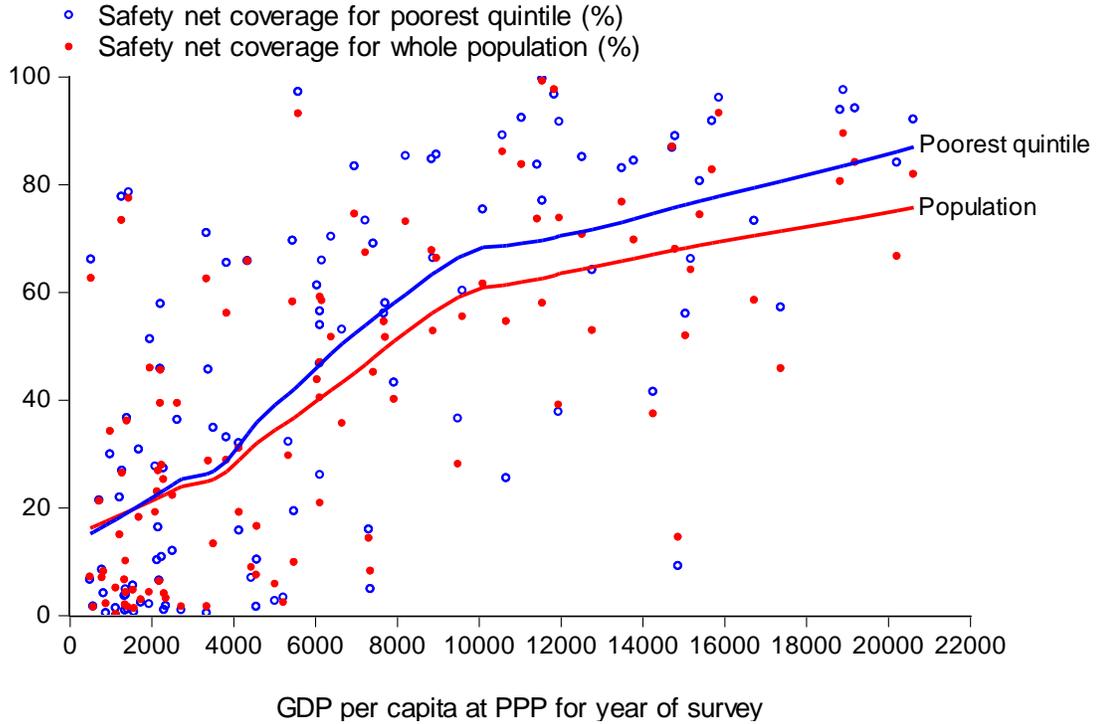


Figure 11: The share of the poorest 20% receiving help from the social safety net in developing countries



Source: Safety net spending includes social insurance and social assistance, including workfare programs. Social safety net coverage rates for poorest quintile (poorest 20% ranked by household income per person) from the World Bank’s ASPIRE site: http://datatopics.worldbank.org/aspire/indicator_glance. The data are available for 109 countries; the latest available year is used when more than one survey is available. GDP from *World Development Indicators*.

Table 1: Estimated consumption floors for the developing world

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Estimated consumption floor ($\hat{E}(x^{\min} y)$)						
	$\alpha = 1$		$\alpha = 2$		Mean consumption	Mean consumption of those living below \$1.25 a day	Contribution of inequality (variance normalized by the mean gap)
	$y^* = \$1.25$		$y^* = \$1.00$	$y^* = \$1.25$			
Sample:	Full	Consumption surveys only	Full	Full	Full	Full	Full
1981	0.59	0.63	0.52	0.51	2.70	0.73	0.14
1984	0.64	0.65	0.56	0.57	2.80	0.80	0.15
1987	0.66	0.66	0.56	0.59	2.97	0.82	0.16
1990	0.67	0.67	0.57	0.60	3.05	0.82	0.15
1993	0.66	0.68	0.56	0.59	3.11	0.83	0.16
1996	0.68	0.69	0.56	0.60	3.18	0.84	0.16
1999	0.68	0.69	0.57	0.61	3.19	0.84	0.16
2002	0.68	0.70	0.57	0.62	3.48	0.85	0.17
2005	0.69	0.71	0.56	0.63	4.03	0.87	0.18
2008	0.69	0.70	0.56	0.63	4.66	0.88	0.18
2010	0.69	0.70	0.54	0.63	5.00	0.88	0.19
2011	0.68	0.69	0.53	0.62	5.23	0.88	0.20

Notes: All numbers are \$ per person per day in 2005 prices using purchasing power parity rates for private consumption. Source: Author's calculations. Columns (1) and (3) use the estimates of PG and SPG from [PovcalNet](#) and equation (7). Column (2) was calculated using *PovcalNet* but only for consumption surveys. Column (4) uses the approximation in equation (12).