Lifting the Floor? Economic Development, Social Protection and the Developing World's Poorest[1]

Juan Margitic and Martin Ravallion
Department of Economics
Georgetown University, Washington DC, 20057, USA

13 October 2018; Revised 1 March, 2019

Abstract: It is theoretically ambiguous whether people in richer countries have a higher floor to their living standards. Nor is it clear whether social protection spending reaches the poorest and thus lifts the floor. Across countries, the paper finds that higher mean incomes come with a higher floor. The bulk of this is direct rather than via public spending on social protection. Social insurance (mainly public pensions) does the "heavy lifting" of the floor. Social assistance (mainly targeted cash-transfers) lifts the floor by only 1.5 cents per day on average, which is less than 10% of mean spending on social assistance.

Keywords: Poverty; inequality; floor; social insurance; social assistance
JEL: I32, I38, O15
Corresponding author: Martin Ravallion, mr1185@georgetown.edu.

[1] For their comments and discussions on this topic, the authors are grateful to Harold Alderman, Abhijit Banerjee, François Bourguignon, Christelle Dumas, Dean Jolliffe, Garance Genicot, Martin Huber, Nishtha Kochhar, Nora Lustig, Karen Macours, Thomas Piketty, François Roubaud, Sylvie Lambert, Dominique van de Walle and seminar participants at the Paris School of Economics, the University of Fribourg, Switzerland, the International Food Policy Research Institute, and Georgetown University. A first draft of this paper was written while the second author was on sabbatical leave at the Paris School of Economics. The authors thank the journal's editor, Andrew Foster, and three anonymous referees for helpful comments.

1. Introduction

Policy discussions often emphasize the need to assure that the poorest are not being "left behind." For example, the title of the 2017 Policy Paper of the UK's Department for International Development (DFID) is "Leaving No One Behind: Our Promise," and the paper's main theme is DFID's goal of prioritizing "the poorest of the poor." One can find many prominent examples of public claims suggesting that DFID's concern is neither isolated nor unjustified—claims that the poorest are in fact being "left behind."[2]

This concern echoes an important school of moral philosophy that has argued that we should judge a society's progress by its ability to enhance the living standards of the poorest, as exemplified by the principle of justice proposed by Rawls (1971). This principle is often advocated for practice, including in the UN's Sustainable Development Goals, which advocate effort to "ensure no one is left behind" (UN, 2017), and the Swiss Constitution, which states that: "…the strength of a people is measured by the well-being of its weakest members."

But are the poorest actually being left behind? This can be interpreted as referring to the lowest level of material living in society. That lower bound can be called the "floor." If the poorest have indeed been left behind then the floor will have stayed put. If the poorest benefit from development and/or social policies then the floor will be lifted. This idea of the floor should not be confused with the "biological floor." Human physiology makes it highly plausible that there is a biological minimum, given that there are strictly positive nutritional requirements for basal metabolism and normal activities. However, economic development and the institutions of (private and public) redistribution can in principle assure that the lower bound is lifted above the biological minimum. The question is whether, and to what extent, that happens in reality.

On *a priori* grounds, it is unclear whether the floor will be lifted by economic development.[3] That will depend on the sources of growth, initial conditions in a country and

---

[2] For example, in 2011, the U.N.'s Secretary-General Ban Ki-moon claimed that: "The poorest of the world are being left behind. We need to reach out and lift them into our lifeboat." (This was at the launch of the report: United Nations, 2011). Similarly, 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." Also, the Vatican's representative to the United Nations claimed in 2015 that the poorest of the world are being left behind (James, 2015).

[3] There is a literature on the effect of economic growth on measures of poverty, including Ravallion (1995), Dollar and Kraay (2002), Ferreira and Ravallion (2009) and Datt et al. (2018); see the survey in Ravallion (2016a, Chapter 8). None of this literature has looked at whether economic growth has benefited the poorest and (hence) raised the floor. The only exception is Ravallion (2016b).

social policies, financed in part from the gains from a higher mean income. For example, growth stemming from external trade expansion could have very different effects on the wages of relatively unskilled workers in poor countries versus richer ones, with different implications for the floor to living standards.[4] Concerns that the poorest may see little or no gain from economic development have prompted social-policy responses. Direct interventions have long been used against poverty in rich countries and are becoming popular in poorer ones. Following the World Bank's usage, we shall label these programs "social protection" (SP).[5] SP coverage in the developing world has expanded rapidly over the last 20 years, with one or more programs now found in most countries (Ravallion, 2016a).[6] However, there are continuing concerns that such efforts are not having much impact on poverty.[7] Various reasons are given including a lack of political will, weak administrative capacity for implementation, ignorance of their rights among poor people, and social stigma associated with targeted programs.

The main approach to assessing the poverty impacts of social policies has compared measures of poverty before and after policy intervention. For transfer policies, this is typically done by comparing measures based on the observed gross income distribution with those obtained by subtracting the transfer received by each household.[8] This can be called the "counting approach." In practice, the aggregate poverty measure is usually a population-weighted average of individual measures.[9] The most widely used measure in practice is the headcount index, giving the proportion of the population living below the poverty line. Higher weight can also be put on poorer people. An example is the Foster-Greer-Thorbecke (FGT)

---

[4] On the diverse distributional implications of trade openness see Ravallion (2006) and Winters et al. (2004).

[5] The policies concerned have had various labels, including "anti-poverty programs," "targeted interventions," "social safety nets," "social assistance," "social insurance" and "social protection." Public spending on SP is sometimes called "social spending" although this term might also be taken to include public health and education spending. To avoid any confusion we will use the more precise term "SP spending."

[6] In population coverage, the two largest SP programs are probably China's *Di Bao* program (a cash program targeted to the poor) and India's *National Rural Employment Guarantee Scheme* (a workfare scheme), both of which can be interpreted as efforts to lift the floor—to assure a minimum standard of living above the biological floor.

[7] With regard to the two examples in the footnote above, on the *Di Bao* program see Ravallion and Chen (2015) and on the India program see Dutta et al. (2014). Evidence on the under-coverage of poor people in cash transfer programs (in Latin America) can be found in Robles et al. (2015). Casual observations of specific antipoverty policies in practice have also expressed concerns about leaving the poorest behind. For example, an article in the *Economist* magazine (2015) on China's poor-area development programs asked how much those programs have helped reduce poverty, and the article's answer referred to how little living standards had risen in one clearly very poor village (in Shanxi) that had apparently been left behind.

[8] This has long been the main approach to benefit-incidence analysis (Kakwani, 1986; van de Walle, 1998). Recent examples include Lindert et al. (2006), Martinez-Vazquez (2008), Lustig et al. (2014) and Piketty et al. (2018).

[9] Atkinson (1987) characterizes this class of measures in more formal terms.

(1984) squared poverty gap (SPG), which weighs poverty gaps by poverty gaps in forming the aggregate measure, thus reflecting inequality among the poor.[10]

The literature on the counting approach has found that absolute poverty measures tend to be lower in countries with a higher mean income, and that these measures tend to fall in growing economies.[11] There is also evidence that social protection spending has generally reduced poverty when measured using the counting approach. For example, in the cross-country data set that we use in this paper we find that SP spending in developing countries roughly halves the average poverty gap index (the agregate gap below the poverty line normalized by the line).[12] The counting approach suggests that, as a rule (and there are exceptions), economic development and social protection tend to reduce poverty.

While the counting approach is of obvious interest and importance, it does not adequately address prevailing concerns about whether the poorest are being left behind. To illustrate the inadequacy of prevailing approaches, Figure 1 shows two pairs of cumulative distribution functions with and without a social protection policy. There is first order dominance in both panels (a) and (b)—an unambiguous change in the aforementioned class of standard (additive) poverty measures. But there is a big difference. In (a), the floor has not risen, but in (b) it has. The poorest in (a) have been "left behind." Existing poverty measures (including those that give higher weight to poorer people) can readily fall without any change in the floor, as in Figure 1, panel (a). Instead, we need to measure the floor, side-by-side with the counting approach.

The following section discusses how social protection and economic development might impact the floor in theory. Section 3 turns to the evidence. We present new evidence on the level of the floor, how much it responds to public spending on SP, and the differences between countries in the efficacy of that spending in lifting the floor. Section 4 concludes.

2. Social protection and the floor in theory

We first explore theoretically how one might expect the floor to respond to both economic development—defined as a rising mean income—and public spending on social

[10] This assures that (unlike the headcount index) a transfer from a poor person to someone even poorer will reduce measured poverty. On this and other axioms for a desirable poverty measure, see Zheng's (1997) survey.
[11] For an overview of the evidence see Ravallion (2016, Chapter 8).
[12] Further evidence on the benefits to poor people as a whole from social spending can be found in Anand and Ravallion (1993), Gupta et al. (2003) and Lustig et al. (2014).

protection. We discuss our measures and data in detail in Section 3 but for now we can simply imagine an observed distribution of household income, y_i, *i*=1,…*n*, and a corresponding (unobserved) distribution y_i^* after eliminating the ignorable transient effects and measurement errors. Let the floor of the y_i^* distribution be denoted $y_{min}^* = \min(y_i^*, i = 1, \dots n)$.

The floor with and without social protection: We can distinguish the pre-transfer floor (denoted y_{min}^{*pre}) from its post-transfer value (y_{min}^{*post}). We are interested in how y_{min}^{*post} varies with both aggregate SP spending per capita, denoted τ, and overall economic development as measured by the mean *m* of the observed distribution. It is assumed that SP spending is financed by domestic taxes, so that *m* is the same before and after transfers.

We can think of y_{min}^{*pre} as being determined by how the overall mean is shared within an economy, while the gain in the floor from SP spending, $y_{min}^{*post} - y_{min}^{*pre}$, is determined by how that spending is shared, which will be the outcome of the processes determining the allocation of SP spending more generally. This suggests a separable structure of the form:

$$y_{min}^{*post} = \vartheta(m) + \varphi(\tau) \quad (1)$$

Here $\vartheta(.)$ and $\varphi(.)$ are the sharing functions determining the pre-transfer floor and the gains from SP respectively. These functions need not be increasing. With regard to $\vartheta(.)$, suppose that economic growth (a higher *m*) is generated in a way that shifts the composition of aggregate domestic demand away from less educated workers. Then the floor could readily fall as the mean rises. Similarly, with regard to $\varphi(\tau)$, there may be general equilibrium effects of SP spending that yield welfare losses to some poor people; this could happen if the poorest benefit little from that spending but end up paying higher prices for the goods they consume.[13] Nor should it be presumed that $\varphi(\tau)$ is linear even if it is increasing. At low levels of spending the gains may well be captured by the non-poor, with the poor only benefiting at higher levels.[14]

A special case is when the poorest receive the mean SP spending, which we later test as the null hypothesis that:

$$y_{min}^{*post} = y_{min}^{*pre} + \tau \quad (2)$$

[13] For example, in the context of a large Philippine cash transfer program, Filmer et al. (2018) find evidence of such adverse welfare effects for non-participating poor people.
[14] For a theoretical model with this property and supportive evidence see Lanjouw and Ravallion (1999).

However, there are reasons why $y_{min}^{*post} - y_{min}^{*pre}$ could differ from τ, including successful efforts at targeting the poorest ($y_{min}^{*post} - y_{min}^{*pre} > \tau$), administrative costs, losses due to corruption, or social exclusion of the poorest (all resulting in $y_{min}^{*post} - y_{min}^{*pre} < \tau$).

Our empirical work follows standard practice in benefit-incidence analysis of estimating the pre-transfer distribution by subtracting transfers received at the household level. This ignores behavioral responses such as through savings, labor supply or private transfers. In defense, it might be argued that strong behavioral responses are unlikely among the poorest, who have the least scope for substitution. However, that might be considered a strong assumption. We provide a partial test of that assumption. The transfer received by the poorest is denoted τ_{min}. If there are behavioral responses by the poorest then $\hat{y}_{min}^{*pre} = y_{min}^{*post} - \tau_{min}$ will underestimate the true value, y_{min}^{*pre}. The extent of the error due to behavioral responses is $b \equiv y_{min}^{*pre} - \hat{y}_{min}^{*pre} \geq 0$. Our test assumes that: (i) the true value of y_{min}^{*pre} is a function of the mean, *m* (as discussed above), and (ii) the behavioral effect *b* is non-decreasing function of mean spending. Thus:

$$\hat{y}_{min}^{*pre} = \vartheta(m) - b(\tau) \quad (b'(\tau) \geq 0) \tag{3}$$

The test is then to see if there is a correlation between the estimated pre-transfer floor and mean spending controlling for mean income. Intuitively, when it is correctly measured, $\hat{y}_{min}^{*pre}$ should not vary with the level of SP spending at a given mean income. Note that, even without behavioral responses by transfer recipients, a negative correlation between $\hat{y}_{min}^{*pre}$ and τ could also arise from a political-economy response of SP transfers to a low level of the floor. So finding such a correlation does not imply the existence of behavioral responses. All we can claim is that finding zero correlation is consistent with our assumption that such responses are absent at the floor. So our test based on (3) can best be thought of as a consistency check on our empirical analysis.

The separablility in (1) might also be considered a strong assumption. A higher mean income may well come with administrative capabilities (including better information systems) that allow governments to better reach the poorest and so raising the marginal gains in the floor from extra social spending. Two examples illustrate how this can happen. First, suppose that economic development brings structural changes such that a rising share of national income is derived from formal-sector activities amenable to taxation. Engels Law implies this as long as the income elasticity of demand for informal sector activities is less than unity. Given that agriculture is the main informal sector in developing countries it is reasonable to assume that economic growth in such countries comes with formalization, generating greater administrative capability including for effective SP. Then it can be expected that economic development allows higher public spending on SP and supports a greater capacity to make that spending effective in reaching the poorest.

The second example starts with the observation that the lack of knowledge about how to access public programs has often been identified as a factor weakening the coverage of poor people by social protection policies.[15] At the same time, economic development tends to come with higher literacy rates, which can be expected to promote greater knowledge, and greater efficacy in dealing with public administrations. Then the marginal gains to the poorest from higher SP spending will tend to rise with mean income when comparing different countries.

In the light of these concerns, instead of equation (1) we write the relationship in the more general form:

$$y_{min}^{*post} = f(\tau, m) \qquad (4)$$

Here f is some (smooth) function (which need not be increasing in its arguments, as noted). So the pre-transfer floor is $y_{min}^{*pre} = \vartheta(m) = f(0, m)$. We shall test separability. When the cross-partial derivative $f_{\tau m}$ is positive we will say that there is weak complementarity. The degree of complementarity plays a role in how economic development impacts the floor, as discussed further below.

It is also of interest to know how much differences in the impact of SP on the floor stem from differences in the overall level of spending versus differences in transfer efficiency. For this purpose we measure what we term Floor Transfer Efficiency (FTE), defined as:

$$FTE = (y_{min}^{*post} - y_{min}^{*pre})/\tau \qquad (5)$$

We also measure the efficiency of transfers in reaching poor people as a whole. Here a standard measure in the literature is what we term Gap Transfer Efficiency (GTE), defined as the reduction in the aggregate poverty gap per $ spent.[16]

Economic development, social protection and the floor: We can identify two channels in how economic development impacts the floor. The first is direct, in that it holds at any given

[15] See Ravallion et al. (2015) for a workfare program in India and Daponte et al. (1999) for food stamps in the US.
[16] GTE is standard output in the ADePT Social Protection software used by the World Bank (Tesliuc and Leite, 2010), although there it is called the "cost-benefit ratio." We prefer our terminology.

level of SP spending. This channel arises through the distribution of the market income gains associated with economic growth. Intuitively, the more "pro-poor" the growth process—such as the more it augments demand for relatively unskilled labor—the stronger is this direct channel. However, being "pro-poor" is not the same thing as reaching the poorest, as discussed in the Introduction; poverty measures can fall yet the poorest are left behind. Indeed, we may see a sinking floor with certain growth processes as noted above.

The second channel is indirect, via higher SP spending. As has long been recognized, a potentially important channel by which economic growth can reduce poverty is via higher SP spending.[17] But is this channel important in practice, and does it embrace the poorest? Economic growth may be heavily concentrated among an elite who use their economic power to further reinforce their positions by promoting political opposition to redistributive tax and spending policies, with implications for the poorest as well as many others. Alternatively, the growth may come with similar or even large gains to electorally influential middle-class citizens who then support anti-poverty efforts, for either altruistic reasons or as insurance given the down-side risks they face. We will be interested in the combined effect of these two channels as well as the components, to see how the level of the floor varies with the level of economic development allowing social policies to adjust.

A simple theoretical model of the political economy of the indirect channel provides some insights. In keeping with Meltzer and Richard (1981), let us assume that the overall level of SP spending is chosen by the median voter. (This can best be interpreted as the median among those who vote, rather than the overall median.) In the present context, what is the relevant distribution for identifying the median voter? Even if $y_i^* - y_i$ has zero mean, the observed median need not equal the true median. One might argue that the observed median is more relevant to the political economy of transfer policy, as this reflects transient factors that could still sway electors. Against this view, the observed distribution also includes measurement errors that may or may not matter to electoral outcomes. Here we will assume that the relevant median, y^{med}, is that of the observed distribution among the electorate. The model can be modified to allow the alternative assumption that it is the median of the y_i^* distribution that matters.

[17] See, for example, the discussion in Anand and Ravallion (1993). The UN's *Human Development Reports* have often emphasized this channel; see, for example, United Nations (2016).

A uniform tax τ is levied to finance SP spending, which depletes the current net income of the median voter.[18] The median voter is taken to face the average tax needed to finance the spending on SP (though this can be relaxed without any important change to the results). The median voter cares about the welfare of poorer people, including those living at the floor. This could be due to altruism or a self-interested concern about down-side vulnerability, including the prospect of personally falling to the floor in the future. And the median voter is assumed to take account of the effects of higher SP spending on the living standards of poorer people. To keep the analysis simple we will ignore levels of living between the floor and the median; this can be readily relaxed but doing so does not appear to deliver any new insights. Thus the median voter cares about utility at the median, net of the tax bill for SP, and utility living at the floor, discounted for the probability of living at the floor in the future. So the choice of τ maximizes:

$$u(y^{med} - \tau) + \rho u(y_{min}^{*post}) \quad (0 < \rho < 1) \tag{6}$$

The utility function, $u(y)$, is taken to be strictly increasing and concave, and ρ is the altruism weight on the utility of the poorest, or the probability of falling to the floor in the future (depending on the interpretation).[19] We allow the possibility that $f_{\tau\tau} > 0$, but that this is bounded above such that:

$$\frac{f_{\tau\tau}}{f_\tau^2} < \frac{-u_{yy}\left(y_{min}^{*post}\right)}{u_y\left(y_{min}^{*post}\right)} \tag{7}$$

When combined with our assumption that $u(y)$ is strictly concave for all y, the condition in (7) assures that the second-order condition for a unique optimal level of SP spending is satisfied.

The median voter's optimal spending on SP, given y^{med} and m, solves:

$$u_y(y^{med} - \tau) = \rho u_y(f(\tau, m)) f_\tau(\tau, m) \tag{8}$$

We can write the solution as:

$$\tau = \tau(y^{med}, m) \tag{9}$$

with first derivatives:[20]

---

[18] Instead one can posit a tax on the median voter that is an increasing function of τ without changing the main argument.

[19] To allow for levels of living between the median and the floor one would add to (6) a sequence of terms in the discounted utility of each of these intermediate incomes, with differing ρ's, and with each income being a (different) function of τ. As noted, this complicates the maths without any obvious extra insight, so we adopt the simpler formulation in (6).

[20] We treat ρ as a constant in the following derivation. Instead, one might prefer to assume that altruism develops as the mean income rises—that altruism gets little weight in very poor societies. Then ρ can be treated as a rising function of m. This adds an extra positive effect to τ_m in the following analysis.

$$\tau_{y^{med}} = \frac{u_{yy}(y^{med}-\tau)}{u_{yy}(y^{med}-\tau)+\rho[u_y\left(y_{min}^{*post}\right)f_{\tau\tau}+u_{yy}\left(y_{min}^{*post}\right)f_\tau^2]} > 0 \qquad (10.1)$$

$$\tau_m = \frac{-\rho[u_y\left(y_{min}^{*post}\right)f_{\tau m}+u_{yy}\left(y_{min}^{*post}\right)f_\tau f_m]}{u_{yy}(y^{med}-\tau)+\rho[u_y\left(y_{min}^{*post}\right)f_{\tau\tau}+u_{yy}\left(y_{min}^{*post}\right)f_\tau^2]} \qquad (10.2)$$

While $\tau_{y^{med}} > 0$ (given (7) and $u_{yy} < 0$), the sign of τ_m is ambiguous. The model allows the possibility that a higher mean at given median—interpretable as higher "inequality"—lowers SP spending. A key issue here is the degree of complementarity between SP spending (higher τ) and economic development (higher *m*) in raising the floor, as indicated by the cross-partial derivative $f_{\tau m}$ (equation 10.2). Complementarity can arise in a number of ways. Countries that are more developed economically may well have greater administrative capabilities for reaching the poorest of the poor. This may also reflect specifics about the type of SP spending; if this facilitates the promotional objective whereby poor people receiving transfers are empowered or incentivized to participate directly in economic development then there is complementarity. Suppose that $f_\tau > 0,\ f_m > 0$ and that:

$$\frac{f_{\tau m}}{f_\tau f_m} > \frac{-u_{yy}\left(y_{min}^{*post}\right)}{u_y\left(y_{min}^{*post}\right)} \qquad (11)$$

If this condition holds then we will say that there is <u>strong complementarity</u> between economic development and SP spending in how they influence the level of the floor. It is evident from (10.2) that strong complementarity implies that $\tau_m > 0$. However, suppose instead that the separability in (1) holds, or that there is substitutability between a higher mean income and SP spending in determining the floor ($f_{\tau m} < 0$). Then we have $\tau_m \leq 0$.

When we consider the bivariate relationship between SP spending and economic development we need to bring in the effect of a higher mean on the median. The total effect of economic development on SP spending is:

$$\frac{d\tau}{dm} = \tau_m + \tau_{y^{med}}\frac{dy^{med}}{dm} \qquad (12)$$

Intuitively, the higher the impact of *m* on y^{med} the more "equitable" the growth process can be said to be. (Indeed, we can think of m/y^{med} as an indicator of inequality as noted.) Of course, the implications for the floor also depend on the incidence of SP spending.

Though our model is simple, it can be used to illustrate a wide range of possibilities. Consider the following stylized, but illustrative, cases.

Case 1: Equitable growth brings both a direct and indirect gain to the poorest. In this case, growth in the mean lifts the floor directly ($f_m > 0$) as well as indirectly via SP spending. Sufficient conditions for the latter channel are that growth in the mean also lifts the median, $\frac{dy^{med}}{dm} \geq 0$, and that there is strong complementarity. Then the effect on the floor is:

$$\frac{dy_{min}^{*post}}{dm} = f_m + f_\tau \frac{d\tau}{dm} > 0 \qquad (13)$$

Recall that if the function $f(\tau, m)$ only exhibits weak complementarity (or substitutability) then the sign of τ_m reverses. It is still possible to find that $\frac{d\tau}{dm} > 0$ and (hence) $\frac{dy_{min}^{*post}}{dm} > 0$; the necessary and sufficient condition for $\frac{d\tau}{dm} > 0$ is that:

$$\rho[u_y(y_{min}^{*post})f_{\tau m} + u_{yy}(y_{min}^{*post})f_\tau f_m] - u_{yy}(y^{med} - \tau)\frac{dy^{med}}{dm} > 0 \qquad (14)$$

Case 2: Inequitable growth leaves the poorest behind. As noted, the direct effect (f_m) could be negative, such as when the specific growth process lowers unskilled wage rates. Here we illustrate a case in which even without a direct effect, the floor falls as the mean rises. This can happen if there is only (at most) weak complementarity and economic development is inequality increasing. Then it is possible to find that neither SP spending nor the level of the floor respond positively to a higher mean income. To illustrate one possible scenario, suppose that $\frac{f_{\tau m}}{f_\tau f_m} \leq \frac{-u_{yy}(y_{min}^{*post})}{u_y(y_{min}^{*post})}$, and that economic growth does not benefit the median voter ($\frac{dy^{med}}{dm} = 0$). Then SP spending falls with a rising mean ($\tau_m < 0$). Furthermore, suppose that the poorest do not share directly in overall economic gains ($f_m = 0$). Then $f_m + f_\tau \tau_m < 0$, i.e., the floor falls as the mean rises. Even without a direct effect, the political economy implies that an inequitable growth process can put a break on the indirect channel, via social spending, thus leaving the poorest behind.

4. Evidence for the developing world

We first describe our empirical measure of the floor. We then implement it for multiple countries in the developing world and expore the patterns in the data, motivated by the theorical model of the previous section. We will show that Case 1 above is more consistent with the data, as a generalization across developing countries.

***Measuring the floor**:* There are limits to how well we could ever hope to measure the floor from standard household surveys. The sampling frame is typically those who live in some form of dwelling, so homeless people and those living in institutions (such as worker dormitories or prisons) are under-represented or even excluded, and they could well be concentrated among the poorest stratum. For example, recent rural migrants in cities living in dormitories or slums could well be under-represented. There has been progress in the design and analysis of multi-frame survey designs that can better represent the homeless.[21] However, practice in using such methods has lagged.

One candidate for the floor is the empirical lower bound of the consumptions measured in a survey. One can think of this as taking the limiting case of the FGT class of measures as the FGT inequality-aversion parameter (α) goes to infinity. However, this would not be satisfactory since there are almost certainly measurement errors and ignorable transient effects in the survey data. For example, all the members of one sampled hosuehold may have been sick during the (often short) recall period used by the survey, and so consumed very little in that period. But one would be loathe to say that they define the floor. There are likely to be such transient effects in the data, whereby observed incomes (or consumption expenditures) in a survey fall temporarily below the floor, but recover soon after. We must also recognize the existence of measurement errors in the cross-sectional survey data available for most countries. Given the measurement errors and transient factors, there is a non-negligible chance that the observed consumption or income of potentially anyone within some stratum of low observed values could in fact be the level of the floor. Some form of averaging is clearly necessary.

Ideally one would use something like the lower bound of time-mean household consumption or income, measured accurately over a much longer period than what is typically measured with survey data. If we were to know the true consumption observed over a long enough period in panel data for a large-enough sample we could reliably estimate the floor directly as this long-run mean. But that is not the data normally available.

Here we follow the approach to measuring the floor in Ravallion (2016b), who estimated the level of the floor for the developign world as a whole, and studed its evolution over time. However, Ravallion (2016b) did not study national values of the floor or the role played by

[21] For example, see Iachan and Dennis (1993). Also see the exposition of these methods in Aenab (2017, Chapter 26).

social protection. Here we do national estimates, and explore how the floor varies with social protection spending and what role a higher mean income plays.

We need an estimator that does not require panel data but can be implemented with cross-sectional surveys, while recognizing the uncertainty as to whether the lowest observed consumption or income in such a survey is in fact the floor. Following Ravallion (2016b) we postulate that any observed income level within a stratum of poor people has some probability of being the floor. These probabilities are not data, but there are some defensible assumptions we can make in lieu of the missing data. While we are uncertain as to whether the lowest observed value is the floor, it is reasonable to assume that this value has the highest probability of being the floor—that our data are sufficiently good to believe that the probability is highest for the person who appears to be the worst off. It also seems reasonable to assume that the probability of being the poorest household declines as the observed measure of income rises. And beyond some point there is no chance of finding the true floor.

We can treat y^*_{min} as a random variable, with a probability distribution given the data. The task is to estimate the mean of that distribution based on the observed incomes:

$$E(y^*_{min}|y) = \sum_{i=1}^{n} \phi_\alpha(y_i) y_i \qquad (15)$$

Here the probability that person *i*, with y_i, is the worst off person is denoted $\phi_\alpha(y_i)$. The specific functional form satisfying these assumptions proposed by Ravallion (2016b) assumes that:

$$\phi_\alpha(y_i) = (1 - \frac{y_i}{z})^\alpha (\frac{1}{nP_\alpha}) \text{ for } 0 \leq y_i \leq z; \alpha \geq 0 \qquad (16.1)$$

$$= 0 \text{ for } y_i > z \qquad (16.2)$$

where $P_\alpha \equiv \frac{1}{n} \sum_{y_i \leq z} (1 - \frac{y_i}{z})^\alpha$ is the FGT class of poverty measures. (The normalization by nP_α assures that the probabilities sum to unity.) For $\alpha > 0$, $\phi_\alpha(y_i)$ attains its maximum value for $y_{min} = \min(y_i \ , i = 1, \ldots n)$ and then falls monotonically with y_i , until it reaches zero at some threshold *z*, above which there is no chance of someone with that income being the poorest.

Under this assumption it can be readily shown that:

$$E(y^*_{min}|y) = z(1 - P_{\alpha+1}/P_\alpha) \qquad (17)$$

Note that the poverty measures are derived from the observed distribution. Thus, the proposed measure of the floor is operational, in that it can be implemented on cross-sectional data.

In the benchmark case we shall set $\alpha = 1$, so that the estimate of the floor is $z(1 - SPG/PG)$ where *SPG* is aforementioned squared poverty gap index ($\alpha = 2$) using *z* as the poverty line and *PG* is the poverty gap index ($\alpha = 1$). However, we will also test robustness to using $\alpha = 0$ and $\alpha = 2$ instead.

Notice that there is nothing to guarantee that a higher mean income lifts the floor, based on this measure. From equation (17), for a fixed *z*, and letting *m* denote the mean of the observed distribution and setting $\alpha = 1$:

$$\frac{\partial \ln E(y^*_{min})}{\partial \ln m} = \left(\frac{\partial \ln PG}{\partial \ln m} - \frac{\partial \ln SPG}{\partial \ln m}\right)\left(\frac{SPG/PG}{1-SPG/PG}\right) \quad (18)$$

Unlike *PG*, the *SPG* index reflects inequality among those living below *z*. It can be seen that lifting the floor requires a larger proportionate reduction in *SPG* than *PG*, i.e., a greater response of the more distribution-sensitive measure. This is a natural consequence of putting higher weight on lower observed incomes when calculating the floor.

Data: Our main data source is the World Bank's "Atlas of Social Protection" (ASPIRE).[22] This draws on household surveys in 122 countries in the developing world, from 1998 to 2014.[23] All currency conversions are at purchasing power parity (for 2005).

In the World Bank's classification, SP spending comprises <u>social insurance</u> (mainly public pension schemes covering old age and disability), <u>social assistance</u> (cash and in-kind transfers and workfare schemes, often targeted to the poor), and <u>labor market programs</u> (training, entrepreneurship support, unemployment benefits). The bulk is social insurance and social assistance. There is clearly a degree of substitutability among these components; if a country is less generous in social insurance it may make up for this using social assistance. We include all components of SP in our analysis. However (as we will see), a large share of SP spending is contributory pensions.[24] We shall comment on the implications of separating out this component as it is rather different given that receipts reflect, in part, past contributions (though governments

[22] ASPIRE is a cross-country compilation of the outputs from a World Bank software program, *ADePT Social Protection*. Tesliuc and Leite (2010) provide a user manual. The ASPIRE team kindly provided detailed output tables from this software by country which we used to build our data set. We assembled the data set in mid-2017.

[23] At the time of writing, 262 surveys are used in the online ASPIRE data set. The special tabulations we requested were only done for the most recent survey for each country. We dropped Zimbabwe from the ASPIRE data as there were clearly serious data quality problems. (There have been numerous problems with Zimbabwe's data in recent times, so this problem was not unexpected.) Whenever SP spending data are used we also dropped Sierra Leone, for which the ASPIRE data show an extremely small positive level of spending relative to the estimated gain in the floor. This may well be a data error. When we take logs the very large negative value for Sierra Leone creates a clear outlier.

[24] Contributory pensions are classified as social insurance by the World Bank; non-contributory social pensions are classified as social assistance.

can still influence current disbursements). We also provide results for social assistance on its own.

One option is to estimate the floor by fixing the poverty line across countries at (say) the World Bank's international line. This approach was rejected as it yields very small subsamples for estimating the floor in many countries, and hence volatile measures. Instead we use poverty lines set at the 0.2 quantile ($q(0.2)$) across all countries, i.e., the poorest 20% in each country define the reference group. The ASPIRE data set provides both *PG* and *SPG*, with and without SP spending. The value of $q(0.2)$ is then held constant for a given country when re-calculating the poverty measures net of transfers. In the ASPIRE dataset, the computations for *SPG* and *PG* pre-transfer are done assuming no behavioral responses. We maintain that assumption, though we provide the test described in Section 3. ASPIRE also provides data on SP transfers received per capita, which we use as our measure of τ.

Summary statistics: Table 1 provides summary statistics. Mean SP spending is $0.88 per person per day. The bulk of this is contributory pensions ($0.67); social assistance accounts for almost all the rest ($0.19). The (un-weighted) mean floor post-transfers is $1.69 a day, though varying widely, from $0.12 to $7.34. There is undoubtedly measurement error; it is very hard to believe that anyone lives at $0.12 per day. While acknowledging the likely measurement errors, we focus on the overall patterns in the data, i.e., the (conditional and unconditional) means. Figure 2 plots the densities of $\hat{y}_{min}^{*post}$ and $\hat{y}_{min}^{*pre}$. The densities are skewed to the right. As we can observe in panel (b) of Figure 2, a log transformation helps to normalize the distributions of both floors. We use this transformation in the bulk of the following analysis. When we study the covariates of the gain in the floor due to SP spending we will use the proportionate gain, $\ln(y_{min}^{*post}/y_{min}^{*pre})$, as our preferred measure.

SP lifted the floor by $0.48 a day on average, as can be seen from Table 1 (comparing post- and pre-transfers). This is well below the mean spending per capita of $0.88. The estimated value of $y_{min}^{*post} - y_{min}^{*pre} - \tau$ is significantly different from zero (*t* = -5.98). Thus, we can reject the null hypothesis in (2). We also observe in Table 1 that SP spending reduced the headcount index by about 7% points on average (recall that the post-transfer index is 20%). There is also a substantial decline in the average poverty gap index, from 10.9% to 5.8%.

The bulk of the impact of SP in developing countries is due to public pensions, which lift the floor by $0.38 a day (Table 1). This too is below the mean spending on such pensions, which is $0.67 per day. Social assistance on its own only raised the floor by $0.015 per day on average—merely 8% of the (already low) level of average spending on social assistance (Table 1). The bulk of the impact of SP on the headcount index (5% points) is also due to contributory pensions. Social assistance on its own reduced the poverty rate by 2% points.

The small impact we find of social assistance in lifting the floor may not be too surprising given the evidence in Brown et al. (2018) (for nine countries in Africa). Their results suggest that prevailing methods of targeting cash transfers based on a "proxy-means test" (PMT) are especially weak in identifying the poorest.[25]

Countries that spend more on social protection tend to have a higher floor. Figure 3 plots the data; the correlation coefficient is 0.751. Mechanically, this relationship reflects both differing levels of SP spending and differing transfer efficiencies. Transfer efficiency in reaching the poorest varies greatly. Figure 4 gives the empirical density function for FTE.[26] (Recall that this is the ratio of the gain in the floor due to SP to mean spending.) We see that very few countries attain a value of FTE of unity or more. For the bulk of countries (87% of the sample), the gain to the poorest is less than mean SP spending. FTE tends to be better for social assistance on its own, for which the median value is 0.934, as compared to a median of 0.630 for all SP; 43% of countries have FTE for social assistance greater than unity.

In addition to FTE, we measure the efficacy of SP in reaching the poorest 20%, giving our second measure of transfer efficiency, GTE. The two measures are correlated, but certainly not perfectly (r = 0.505). GTE is positively correlated with spending per capita (r = 0.656), but that is not true for FTE (r = -0.021). As countries spend more on social protection, a larger share of that spending tends to reach the poorest 20% but not the poorest. Figure 5 plots the relationships with average SP spending for both FTE and GTE (it is easier to see if one logs spending per capita). This points to a notable difference in efficacy in reaching the poorest quintile versus the poorest households. By implication, relative efficiency in reaching the poorest (FTE/GTE) declines with mean spending (r = -0.430).

However, the bulk of the variance in the impact of higher SP spending on the floor (as evident in Figure 3) is due to the variance in aggregate levels of that spending, rather than its

---

[25] An important reason is the widespread use of standard (OLS) regression for calibrating the PMT score (regressing observed consumption on covariates used to predict consumption in the population). The residuals from such a regression will be positively correlated with the dependent variable (with a covariance equal to the variance of the error). Thus, the method will overestimate living standards for the poorest.

[26] Recall that Sierra Leone is dropped; this makes the bulk of the density function easier to see in Figure 3.

efficiency in reaching the poorest. If one decomposes the variance in $\ln(y_{min}^{*post} - y_{min}^{*pre})$ into the variance in log spending per capita, the variance in the log of FTE, and the covariance, the first component accounts for 77%, with the variance in log FTE accounting for 14% and the covariance representing 9%. (Recall that FTE has a low correlation with sending per capita.)

Richer countries tend to have a higher floor. Figure 6 plots the (log) floor, both pre- and post-transfer, against the (log) mean. Also notice that the two regression lines diverge. The pre-transfer floor has an elasticity of about 0.8 to the mean, while it is 0.9 for the post-transfer floor, and the difference is statistically significant (at the 1% level).[27] The income elasticity of the pre-transfer floor is significantly less than unity (t = 3.2), implying that the (pre-transfer) floor tends to fall as a share of the mean as the latter rises. By contrast, the income elasticity of the post-transfer floor is not significantly different from unity, implying that the floor does not fall relative to the mean as economies develop. Thus we see that, on average, SP spending in developing countries is able to negate the tendency for the pre-transfer floor to fall as a share of the mean as the mean rises with economic development.

These elasticities also imply substantial <u>absolute</u> divergence between the floor and the mean. At the mean points from Table 1, a \$1.00 increase in mean income comes with a \$0.11 increase in the pre-transfer floor, and a \$0.19 increase in the post-transfer floor.

Despite this strong correlation between the floor and mean income, the FGT poverty measures do not provide reliable indicators of the level of the floor. Indeed, the (post-transfer) PG and SPG measures have only weak negative correlations with the (post-transfer) floor; r = -0.156 and -0.131 respectively. A much better indicator of the floor is the quantile of the poorest 1% ($q(0.01)$) for which r = 0.825. However, for the reasons noted in Section 2, $q(0.01)$ could be a noisy measure. Our measure can be calculated from the same primary data.

In terms of the model in Section 2, the strong positive relationship between the level of the floor and mean income reflects both higher SP spending in richer countries and a direct effect at given spending. We will now use regressions to separate out these effects.

Partial effects: To allow for multiple covariates, we now explore these relationships further using regressions. (We do not intend that these regressions be given a casual

---

[27] There is a suggestion in Figure 6 that the elasticity is not constant, but declines as the mean rises. Adding a squared term in the log mean, its coefficient is negative and significant at the 5% level in both cases. However, this effect vanishes if one drops the two observations with highest mean.

interpretation, but only as a convenient means of testing for partial correlations and a means of testing our measurement method.)

Recall that an implication of our assumption that the pre-transfer floor is the post-transfer value less SP spending received by the poorest is that we should not find a correlation between the estimated pre-transfer floor and mean spending (Section 2). While there is a significantly positive (zero-order) correlation between the pre-transfer floor and average SP spending (r = 0.511), this vanishes when we control for the mean. The partial correlation falls to 0.068; Figure 7 plots the two series (with log floor predicted at mean income) while Table 2 gives the regression where we also see clearly that countries with a higher overall mean have a higher pre-transfer floor. The restriction that SP spending does not affect the pre-transfer floor at a given mean performs well. This provides support for our estimation method ignoring any behavioral responses of the poorest.

Table 2 also provides the regressions for the gain in the floor attributed to SP, i.e., our estimate of $\ln(y_{min}^{*post}/y_{min}^{*pre})$. We see that higher aggregate transfers contribute to a larger impact of transfers on the floor. Noting that we can obtain the regression for the post-transfer floor by adding that for $\ln(y_{min}^{*post}/y_{min}^{*pre})$ to that for $\ln(y_{min}^{*pre})$, we see that there is both a direct effect of higher mean income on the post-transfer floor and an indirect effect, via higher SP spending; the direct effect is 0.642 (s.e.=0.070). However, when normalized by the total income elasticity of 0.923 (Figure 6), we see that the bulk (70%=0.642/0.923) of the effect is direct.

We find that there is a positive interaction effect between average transfers and the mean, which helps in raising the impact of SP on the floor (Table 2, Column 5). This suggests weak complementarity between SP spending and economic development in how they influence the efficacy of SP in raising the floor; a higher mean income comes with higher marginal gains to the poorest from higher public spending on SP.

One clue to the role played by heterogeneity in transfer effectiveness is to augment the regressions with gap transfer efficiency; recall that this is the impact of SP spending on the aggregate poverty gap for the poorest 20% per $ of spending. Here we are interested in seeing whether countries that are more efficient at reaching the poorest 20% also tend to do better at lifting the floor, and here we can expect both an additive effect and an interaction effect with mean transfers. We can go further and allow a complete set of interaction effects, including with the mean. This augmented specification is in Column (6) of Table 2.

As expected, there is a strong interaction between GTE and transfer spending in their effects on the extent to which SP lifts the floor. There is also a negative interaction effect between mean income and transfer efficiency; it is in poor countries where the effectiveness in transferring money to the poorest 20% tends to matter more to lifting the floor. When we evaluate the total effects at the mean points, we find a significant positive effect of SP spending and GTE on the extent to which those transfers succeed in raising the floor (Table 2, lower panel). Once we control for the level of transfers and transfer efficiency we do not find that higher average incomes come with a greater impact of transfers on the floor.

Robustness tests: We provide two tests. The first relates to our choice of the parameter α, while the second relates to a potential omitted variable.

Recall that our benchmark case assumes that $\alpha = 1$, meaning that the weight on each observed income declines linearly in measuring the floor reaching zero at *z*. To test robustness to this choice, Table 3 provides the analogous results to Table 2 for using $\alpha = 0$ and $\alpha = 2$.[28] In the former case, all income below *z* are weighted equally in measuring the floor (which then becomes mean income below *z*), while in the latter case the weights decline as a quadratic function. The main qualitative results from Table 2 are quite robust to these changes. The main exception is for $\alpha = 0$ for which our consistency test (that the estimated pre-transfer floor should be uncorrelated with mean transfers at given mean consumption) no longer passes. This could reflect incentive effects nearer *z*, given that those income levels are weighted less when using $\alpha > 0$. As we found for $\alpha = 1$, the test passes for $\alpha = 2$.

Turning to the second robustness test, a potentially important omitted variable is the age-dependency ratio (ADR). Countries with a higher ADR can be expected to both spend more on SP, have a lower average income, and a lower floor. Thus an omitted variable bias may be present in our regressions in Table 2. Following common practice, the ADR is defined as the number of people older than 65 years or younger than 14, divided by the number of people between 15 and 64 years of age. We provide in Table 4 the same set of regressions including a control for ADR (for $\alpha = 1$). We find no significant effect of ADR on the pre-transfer floor, but we do find such effects on the post-transfer floor, with a higher dependency ratio implying less

[28] Note that this requires the cubed poverty gap (P_3) but that is not in the ASPIRE meta-data. To estimate P_3 we have used the approximation formula derived in Ravallion (2016b) by taking a second-order Talylor series expansion around the mean income of the poor.

gain in the floor due to social spending. The effect of the latter is slightly lower with the extra control, and the (absolute value of) mean income effect is higher (Table 4).

4. Conclusions

In our theoretical model, the level of the floor—the lower bound of the distribution of living standards in society—depends on both public social protection spending and mean income. The level of SP spending depends in turn on the mean. Thus, there is both a direct effect of economic development on the floor, and an indirect effect via social protection. A key role is played by the extent of complementarity between SP spending and development; complementarity exists when higher SP spending increases the marginal gains to the poorest from a higher mean income. If this complementarity is not too weak, and the growth process is not too inequitable, then the floor will rise with economic development; the poorest will not be left behind. But our model illustrates that there is no guarantee that the poorest will see any gain from overall economic development. That is an empirical issue.

We have provided the first evidence for developing countries. To test whether public spending on social protection has reached the poorest we must be able to measure the floor. This poses a difficult measurement problem. We must recognize that the lowest observed consumption or income in a cross-sectional survey need not be a good indicator of the true floor to living standards given predictable transient effects and measurement errors in the data. To help address these concerns we have followed Ravallion (2016b) in measuring the floor as a weighted mean for those sampled households in a reference group that is assumed to include the poorest. While we cannot be certain that the lowest measured consumption or income is that of the poorest, we assume that the probability is highest at that point, but then declines linearly up to the 20th percentile after which it is assumed that there is no chance of being the floor. Our main findings are robust to relaxing linearity to allow either no decline in the probability or a nonlinear (quadratic) decline.

We have also studied how the impact of SP spending on the poorest varies with both mean income and the level of SP spending. We find that higher SP spending has lifted the floor in developing countries as a whole. The poorest benefit from this spending. There is considerable variability across countries, the bulk of which (in terms of variance) is due to differences in the level of SP spending rather than the transfer efficiency of that spending.

However, for the cross-country sample, the gain from SP for those living at the floor is significantly less than aggregate SP spending per capita. Social insurance (mainly public pensions for old age and disability) does the "heavy lifting" of the floor. Social assistance (mainly cash transfers often targeted to poor people) on its own lifts the floor by merely 1.5 cents per day on average—less than 10% of the (seemingly low) level of mean spending on social assistance. This component of SP is largely missing the poorest in the developing world.

We find that higher average income tends to come with a higher pre-transfer floor, though not enough to prevent a relative decline in the floor as the mean rises, and large absolute divergence. This is the sense in which it can be said that the poorest tend to be "left behind" with economic development. The bulk of the efficacy of a higher mean income in lifting the floor appears to be direct, rather than via higher SP spending. Nonetheless, SP spending comes close to assuring that the post-transfer floor does not sink relative to the mean when comparing low and middle-income countries. Statistically, while the pre-transfer floor tends to fall relative to the mean as the latter rises, we cannot reject the null hypothesis that the post-transfer floor stays at a constant share of the mean. There is also evidence of complementarity between social spending and economic development, as evident in a strong positive interaction effect between SP spending and mean income in regressions for the gains to the poorest from higher SP spending. Along with rising SP spending, this complementarity plays a positive role in helping to assure that the poorest benefit from economic growth.

Our results also reveal the inadequacy of prevailing poverty measures in addressing widespread concerns about the level of living of the poorest. Standard poverty measures show only low (negative) correlation coefficients with our estimates of the floor. Thus, the paper has underlined the importance of focusing on the floor directly. For this purpose, we have put new measures on the table, though we expect that they can be improved, especially with better data on the poorest. While data will hopefully improve, we suggest that an estimate of the floor should now become a staple in the dashboard of social indicators. Doing so would make poverty measurement more relevant to the ongoing concerns among policy makers and citizens about not leaving the poorest behind.

References

Aenab, Raghunath, 2017, *Survey Sampling Theory and Applications*, London: Academic Press.

Anand, Sudhir, and Martin Ravallion, 1993, "Human Development in Poor Countries: On the Role of Private Incomes and Public Services," *Journal of Economic Perspectives* 7: 133-150.

Atkinson, Anthony B., 1987, "On the Measurement of Poverty," *Econometrica* 55: 749-764.

Brown, Caitlin, Martin Ravallion and Dominique van de Walle, 2018, "A Poor Means Test? Econometric Targeting in Africa," *Journal of Development Economics* 134: 109-124.

Daponte, Beth, Seth Sanders, Lowell Taylor, 1999, "Why do Low-Income Households Not Use Food Stamps? Evidence from an Experiment," *Journal of Human Resources* 34(3): 612–628.

Datt, Gaurav, Martin Ravallion, and Rinku Murgai, 2018, "Poverty and Growth in India over Six Decades," Development Policy Working Paper, World Bank.

Dollar, David, and Kraay, Aart, 2002, "Growth is Good for the Poor," *Journal of Economic Growth* 7(3): 195-225.

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, Washington DC.

Economist, The, 2015, "Just a Little Bit Richer," *Economist Magazine*, April 4.

Ferreira, Francisco, and Martin Ravallion, 2009, "Poverty and Inequality: The Global Context," in *The Oxford Handbook of Economic Inequality*, edited by Wiemer Salverda, Brian Nolan and Tim Smeeding, Oxford: Oxford University Press.

Filmer, Deon, Jed Friedman, Eeshani Kandpal, and Junko Onishi, 2018, "General Equilibrium Effects of Targeted Cash Transfers: Nutrition Impacts on Non-Beneficiary Children," Policy Research Working Paper 8377, World Bank.

Foster, James, J. Greer, and Erik Thorbecke, 1984, "A Class of Decomposable Poverty Measures," *Econometrica* 52: 761-765.

Gupta, Sanjeev, Marijn Verhoeven, and Erwin R. Tiongson, 2003, "Public Spending on Health Care and the Poor," *Health economics* 12, no. 8 (2003): 685-696.

Iachan, Ronaldo, and Michael L. Dennis, 1993, "A Multiple Frame Approach to Sampling the Homeless and Transient Population," *Journal of Official Statistics* 9(4): 747-764.

James, Aaron, 2015, "World's poorest left behind, says Vatican," *Christian Radio* 11 February.

Kakwani, Nanak, 1986, *Analyzing Redistribution Policies: A Study using Australian Data.* Cambridge: Cambridge University Press.

Korinek, Anton, Johan Mistiaen and Martin Ravallion, 2007, "An Econometric Method of Correcting for Unit Nonresponse Bias in Surveys," *Journal of Econometrics* 136: 213-235.

Lanjouw, Peter, and Martin Ravallion, 1999, "Benefit Incidence and the Timing of Program Capture" *World Bank Economic Review* 13(2): 257-274.

Lindert, Kathy, Emmanuel Skoufias and Joseph Shapiro, 2006, *Redistributing Income to the Poor and the Rich: Public Transfers in Latin America and the Caribbean*, World Bank, Washington DC.

Lustig, Nora, Carola Pessino, and John Scott, 2014, "The Impact of Taxes and Social Spending on Inequality and Poverty in Argentina, Bolivia, Brazil, Mexico, Peru and Uruguay," *Public Finance Review* 42(3).

Martinez-Vazquez, Jorge, 2008, ''The Impact of Budgets on the Poor: Tax and Expenditure Benefit Incidence Analysis.''In *Public Finance for Poverty Reduction Concepts and Case Studies from Africa and Latin America*, edited by Blanca Moreno-Dodson and Quentin Wodon, Chapter 5. Washington, DC: The World Bank.

Meltzer, Alan, and Scott F. Richard, 1981, "A Rational Theory of the Size of Government," *Journal of Political Economy* 89(5): 914-927.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman, 2018, "Distributional National Accounts: Methods and Estimates for the United States," *Quarterly Journal of Economics* 133(2): 553-610.

Ravallion, Martin, 1995, "Growth and Poverty: Evidence for Developing Countries in the 1980s," *Economics Letters* 48, 411-417.

______________, 2006, "Looking Beyond Averages in the Trade and Poverty Debate," *World Development* 34(8): 1374-1392.

______________, 2016a, *The Economics of Poverty: History, Measurement and Policy*. New York: Oxford University Press.

_______________, 2016b, "Are the World's Poorest Being Left Behind?," *Journal of Economic Growth* 21(2): 139–164.

Ravallion, Martin, Dominique van de Walle, Rinku Murgai and Puja Dutta, 2015, "Empowering Poor People through Public Information? Lessons from a Movie in Rural India," *Journal of Public Economics* 132: 13-22.

Ravallion, Martin, and Shaohua Chen, 2015, "Benefit Incidence with Incentive Effects, Measurement Errors and Latent Heterogeneity: A Case Study for China," *Journal of Public Economics* 128: 124-32.

Rawls, John, 1971, *A Theory of Justice,* Cambridge MA: Harvard University Press.

Robles, Marcos, Marcela Rubio, and Marco Stampini, 2015, "Have Cash Transfers Succeeded in Reaching the Poor in Latin America and the Caribbean?" Policy Brief IDB-PB-246, Inter-American Development Bank, Washington DC.

Tesliuc, Emil D. and Phillippe G. Leite, 2010, *User Manual for AdePT Social Protection*, World Bank, Washington DC.

United Nations, 2011, *The Millennium Development Goals Report,* New York: United Nations Secretariat.

_____________, 2016, *Human Development Report,* New York: UNDP.

_____________, 2017, *The Sustainable Development Goals Report*, New York: United Nations.

van de Walle, Dominique, 1998, "Assessing the Welfare Impacts of Public Spending," *World Development* 26(3): 365-379.

Winters, L. Alan, Neil McCulloch and Andrew McKay, 2004, "Trade Liberalization and Poverty: The Evidence So Far," *Journal of Economic Literature* 42(March): 72-115.

World Bank, 2014, *The State of Social Safety Nets 2014,* Washington DC: World Bank.

Zheng, Buhong, 1997, "Aggregate Poverty Measures," *Journal of Economic Surveys* 11(2): 123-140.

Figure 1: Both pairs of distributions show first-order dominance but with very different implications for the floor

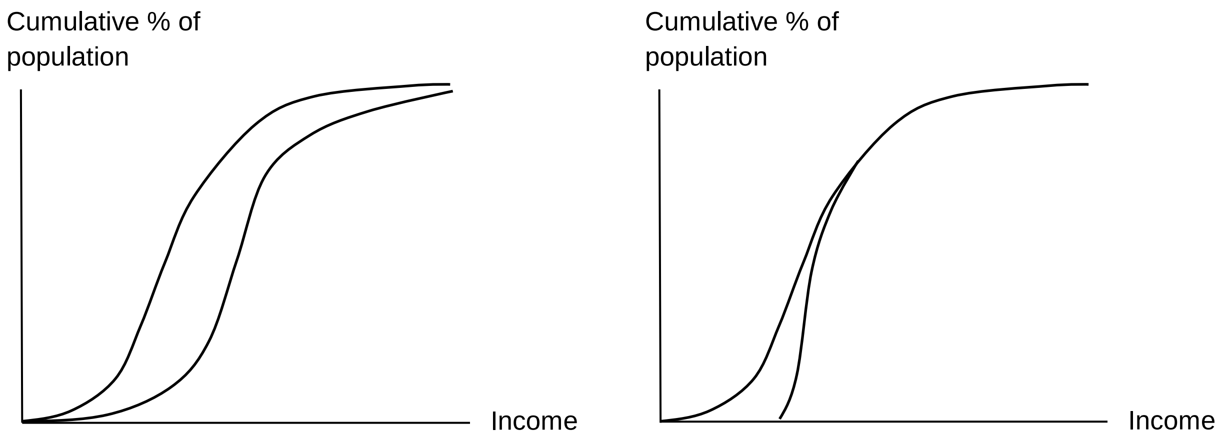

(a) The floor has not risen; poorest are left behind

(b) Floor has risen

Source: Ravallion (2016a).

Figure 2: Kernel density functions for the floor across countries

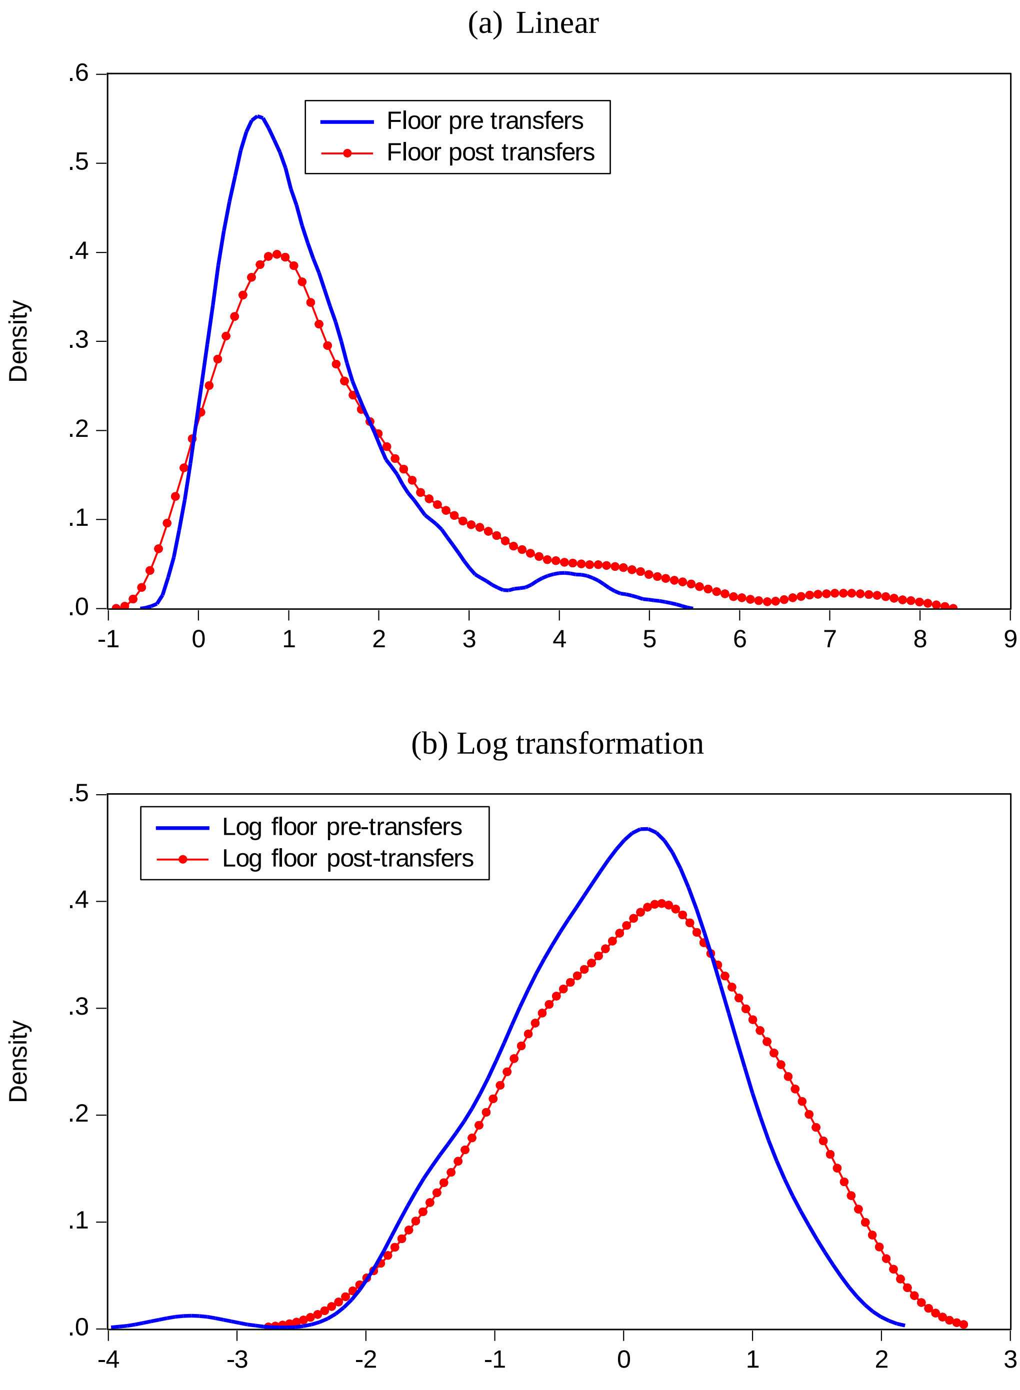

Figure 3: Higher SP spending comes with a higher floor

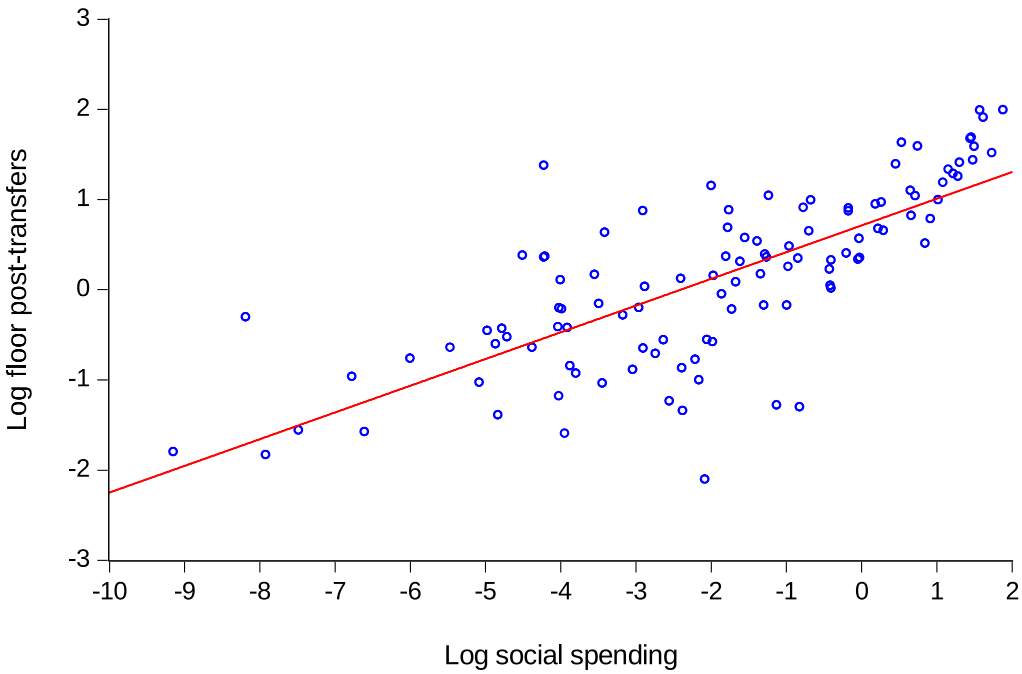

Figure 4: Kernel density functions for floor transfer efficiency

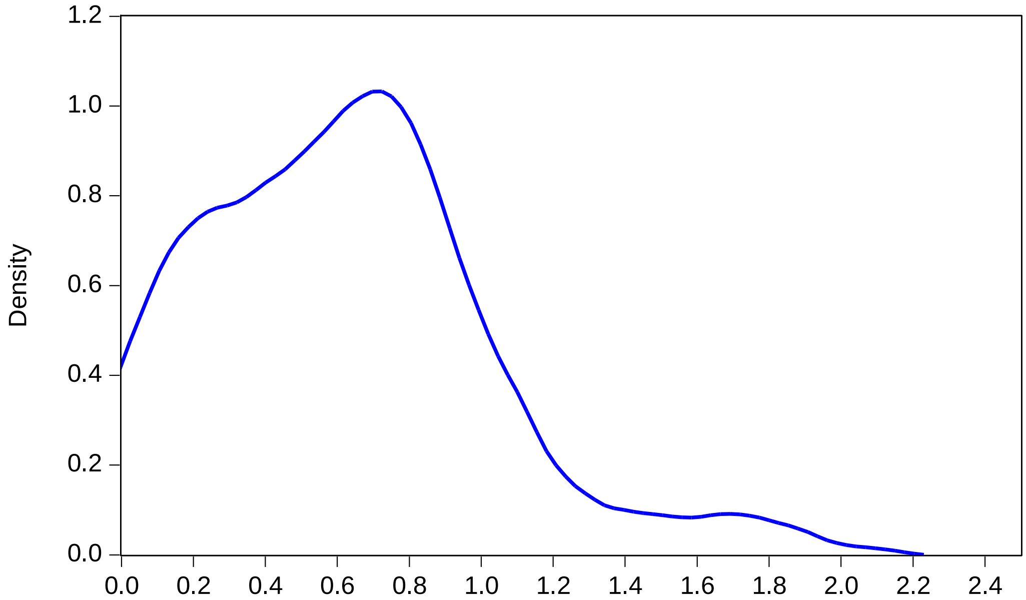

Figure 5: Transfer efficiency plotted against aggregate transfers per capita

(a) Gap transfer efficiency

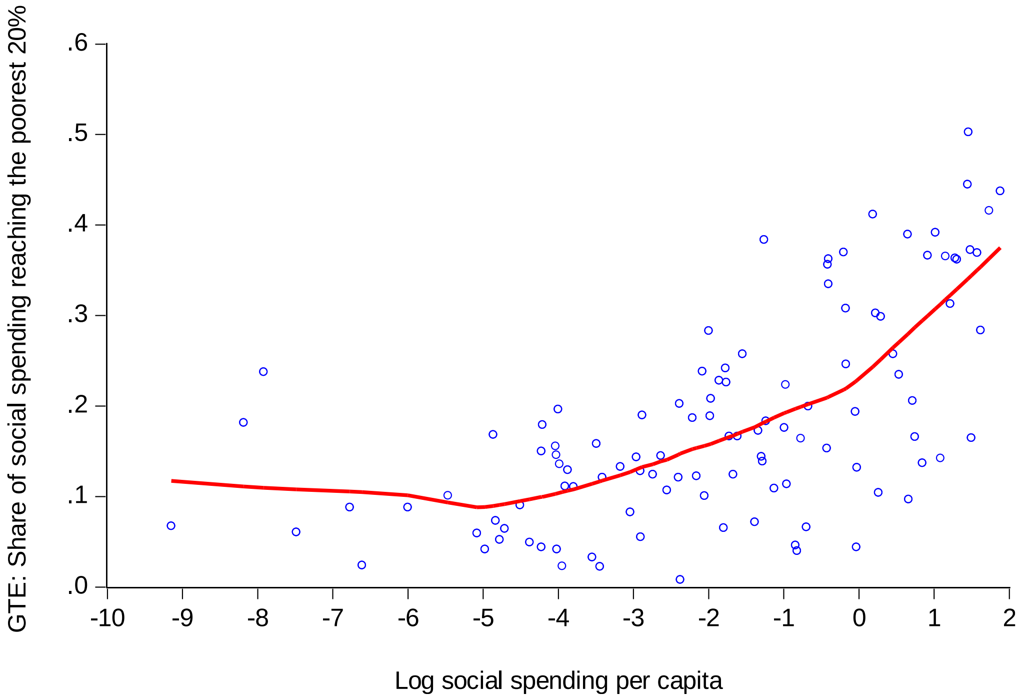

(b) Floor transfer efficiency

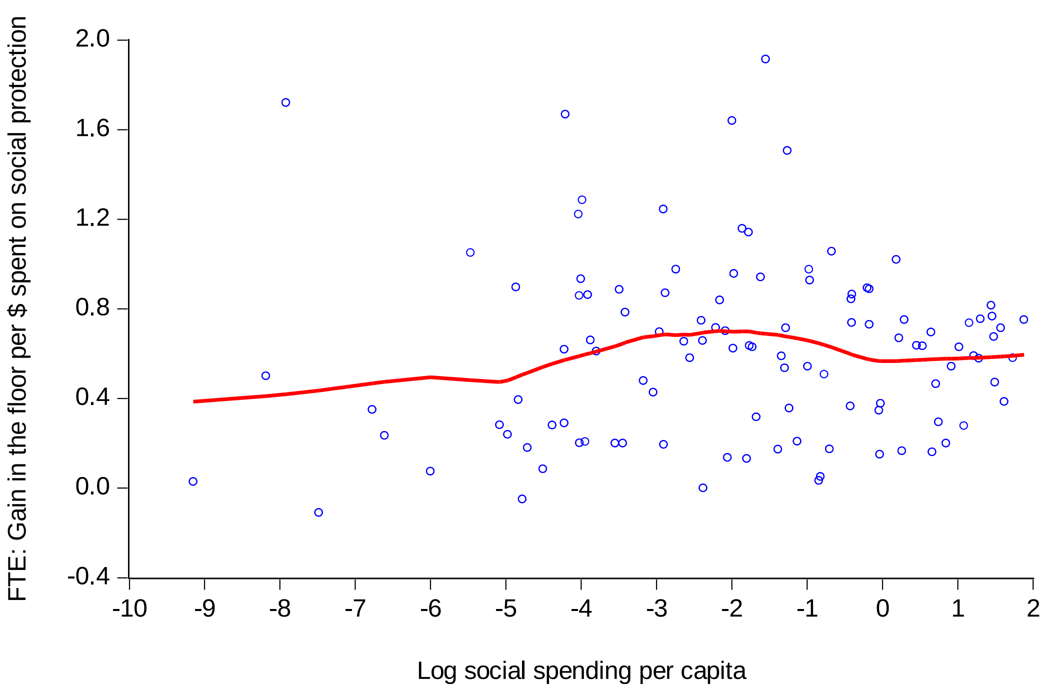

Figure 6: Richer countries have a higher floor

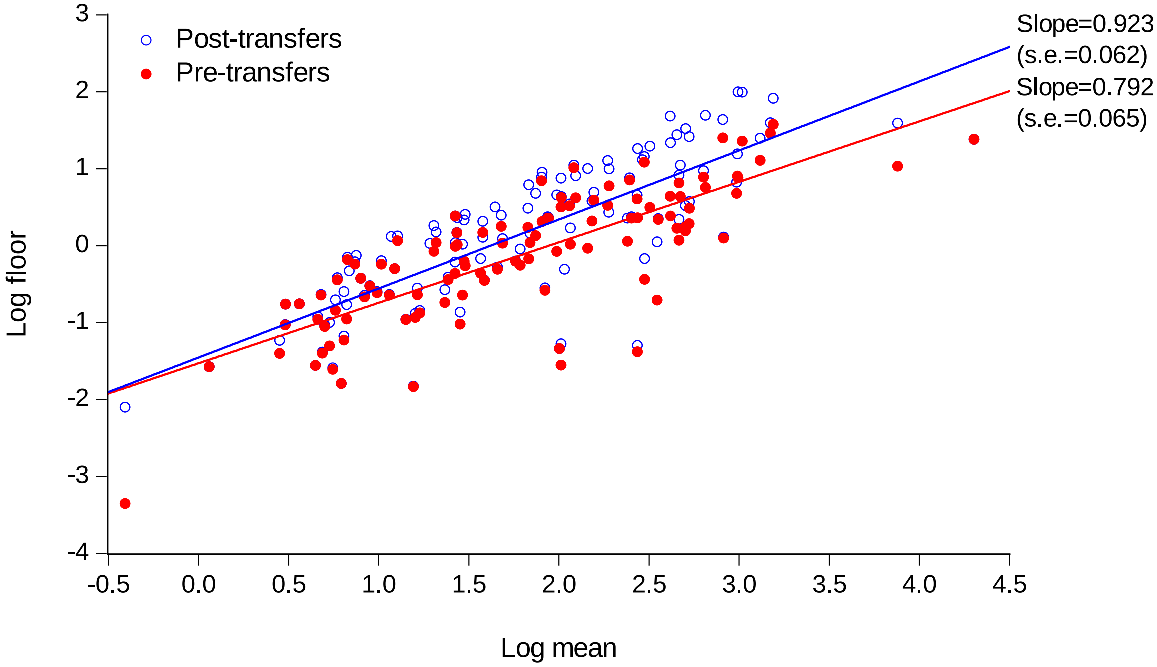

Figure 7: Log pre-transfer floor controlling for the mean plotted against log SP spending per capita

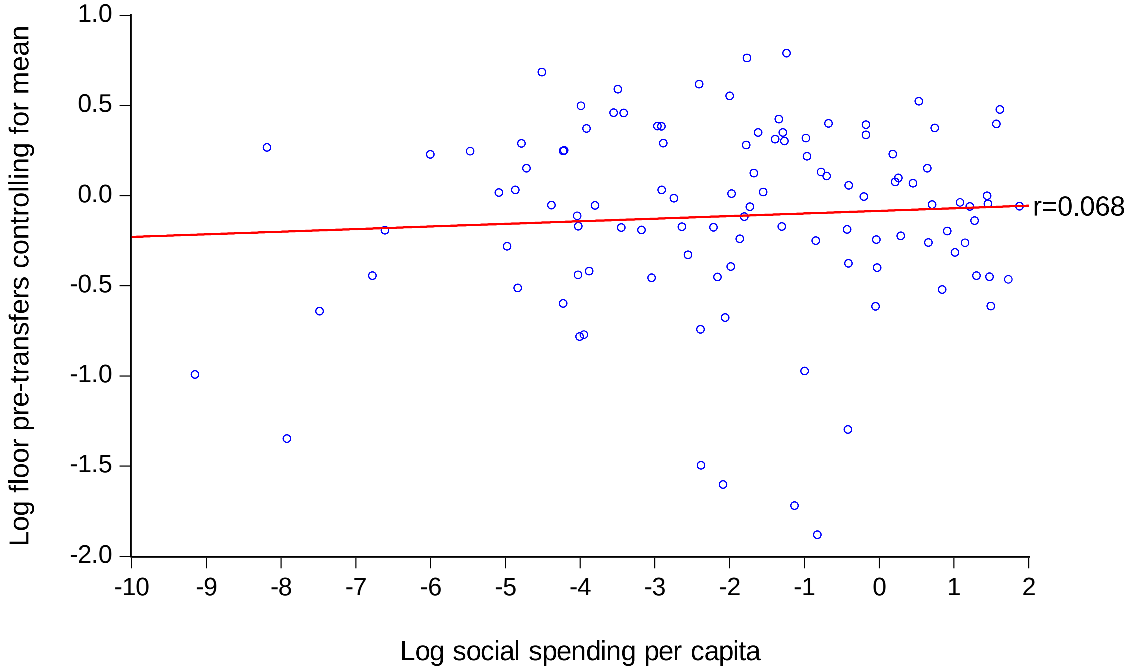

Table 1: Summary statistics

	N	Mean	St. dev.	Median	Min	Max
Survey mean (m)	121	8.413	9.033	5.973	0.67	74.05
Threshold (z)	121	2.979	2.895	1.963	0.20	17.48
Mean social protection spending (τ)	111	0.876	1.437	0.171	0.00	6.56
Mean contributory pensions	116	0.667	1.195	0.109	0.00	5.51
Mean social assistance	93	0.187	0.321	0.062	0.00	1.92
Floor post transfers ($\hat{y}_{min}^{*post}$)	121	1.693	1.547	1.184	0.12	7.34
Floor post transfers as share of threshold	121	0.580	0.095	0.594	0.228	0.729
Floor pre all transfers ($\hat{y}_{min}^{*pre}$)	111	1.210	0.954	1.009	0.03	4.82
Floor pre transfers as share of threshold	111	0.463	0.139	0.487	0.174	0.728
Floor pre contributory pensions only	116	1.308	1.010	1.101	0.11	5.70
Floor pre social assistance only	93	1.678	1.442	1.184	0.04	6.42
Headcount index pre all transfers (%)	111	26.802	9.531	22.390	20.00	56.83
Headcount index pre contributory pensions (%)	116	24.612	7.586	20.830	18.96	49.25
Headcount index pre social assistance alone (%)	93	22.236	3.913	21.000	19.80	49.00
Poverty gap index post transfers (%)	121	5.744	1.547	5.531	3.51	11.19
Poverty gap index pre transfers (%)	111	10.813	7.251	7.780	3.56	36.90
Squared poverty gap index post transfers (x100)	121	2.556	1.384	2.254	0.952	8.594
Squared poverty gap index pre transfers (x100)	111	6.680	6.246	4.011	0.967	30.489

Note: All values displayed above are in daily per capita US$ units, in 2005 prices (at PPP) unless noted otherwise. SP spending comprises all social insurance, social assistance and labor market programs (see text). The number of countries can vary depending on data availability, as indicated.

Table 2: Regressions for log floor ($\alpha = 1$)

	(1)	(2)	(3)	(4)	(5)	(6)
	Log floor, pre-transfers ($\ln(y^{*pre}_{min})$)		Gain in the floor due to SP spending ($\ln(y^{*post}_{min} / y^{*pre}_{min})$)			
Log SP transfers per capita ($\ln \tau$)	0.027		0.118***	0.095***	0.058**	0.244***
	(0.033)		(0.018)	(0.012)	(0.027)	(0.043)
Log mean income ($\ln m$)	0.740***	0.792***	-0.099*		-0.030	-0.347**
	(0.105)	(0.065)	(0.052)		(0.066)	(0.155)
Interaction effect ($\ln \tau . \ln m$)					0.035**	0.025*
					(0.014)	(0.013)
Log Gap Transfer Efficiency ($\ln GTE$)						0.664***
						(0.138)
Interaction effect ($\ln \tau . \ln GTE$)						0.104***
						(0.014)
Interaction effect ($\ln m . \ln GTE$)						-0.151**
						(0.064)
Constant	-1.399***	-1.544***	0.661***	0.439***	0.497***	1.701***
	(0.249)	(0.131)	(0.128)	(0.039)	(0.152)	(0.340)
R^2	0.629	0.626	0.495	0.462	0.536	0.800
Total effect evaluated at mean points						
$\ln \tau$					0.121***	0.088***
					(0.020)	(0.016)
$\ln m$					-0.098	-0.101**
					(0.064)	(0.051)
$\ln GTE$						0.191***
						(0.029)

Note: OLS regressions. Robust standard errors in parentheses. N=110. ***: 1% significance;**: 5%; *10%.

Table 3: Regressions for log floor for alternative values of α

(a) $\alpha = 0$

	(1)	(2)	(3)	(4)	(5)	(6)
	Log floor, pre-transfers ($\ln(y^{*pre}_{min})$)		Gain in the floor due to SP spending ($\ln(y^{*post}_{min}/y^{*pre}_{min})$)			
Log SP transfers per capita ($\ln \tau$)	0.062**		0.064***	0.053***	0.026*	0.133***
	(0.026)		(0.010)	(0.007)	(0.015)	(0.024)
Log mean income ($\ln m$)	0.757***	0.877***	-0.044		-0.001	-0.175**
	(0.082)	(0.052)	(0.027)		(0.035)	(0.081)
Interaction effect ($\ln \tau . \ln m$)					0.022***	0.016**
					(0.008)	(0.007)
Log Gap Transfer Efficiency ($\ln GTE$)						0.367***
						(0.074)
Interaction effect ($\ln \tau . \ln GTE$)						0.059***
						(0.008)
Interaction effect ($\ln m . \ln GTE$)						-0.082**
						(0.034)
Constant	-1.023***	-1.360***	0.340***	0.240***	0.237***	0.904***
	(0.192)	(0.102)	(0.068)	(0.022)	(0.082)	(0.179)
Observations	110	110	110	110	110	110
R-squared	0.732	0.716	0.487	0.466	0.539	0.800

Note: Robust standard errors in parentheses; n=110. *** p<0.01, ** p<0.05, * p<0.1

(b) $\alpha = 2$

	(1)	(2)	(3)	(4)	(5)	(6)
	Log floor, pre-transfers ($\ln(y^{*pre}_{min})$)		Gain in the floor due to SP spending ($\ln(y^{*post}_{min}/y^{*pre}_{min})$)			
Log SP transfers per capita ($\ln \tau$)	-0.008		0.149***	0.138***	0.021	0.316***
	(0.044)		(0.023)	(0.020)	(0.036)	(0.067)
Log mean income ($\ln m$)	0.720***	0.706***	-0.047		0.123	-0.180
	(0.148)	(0.089)	(0.041)		(0.082)	(0.123)
Interaction effect ($\ln \tau . \ln m$)					0.073***	0.030*
					(0.022)	(0.017)
Log Gap Transfer Efficiency ($\ln GTE$)						0.723***
						(0.112)
Interaction effect ($\ln \tau . \ln GTE$)						0.138***
						(0.020)
Interaction effect ($\ln m . \ln GTE$)						-0.107**
						(0.050)
Constant	-1.630***	-1.589***	0.716***	0.611***	0.319*	1.732***
	(0.353)	(0.178)	(0.105)	(0.065)	(0.166)	(0.269)
Observations	110	110	103	103	103	103
R-squared	0.479	0.478	0.479	0.476	0.557	0.800

Note: Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table 4: Regressions for log floor controlling for age-dependency ratio

	(1)	(2)	(3)	(4)	(5)	(6)
	Log floor, pre-transfers ($\ln(y_{min}^{*pre})$)		Gain in the floor due to SP spending ($\ln(y_{min}^{*post}/y_{min}^{*pre})$)			
Log SP transfers per capita ($\ln \tau$)	0.006		0.090***	0.962***	0.029	0.227***
	(0.046)		(0.018)	(0.161)	(0.023)	(0.042)
Log mean income ($\ln m$)	0.718***	0.723***	-0.124**		-0.052	-0.315*
	(0.118)	(0.104)	(0.059)		(0.071)	(0.172)
Interaction effect ($\ln \tau . \ln m$)					0.036**	0.026*
					(0.014)	(0.013)
Log Gap Transfer Efficiency ($\ln GTE$)						0.655***
						(0.149)
Interaction effect ($\ln \tau . \ln GTE$)						0.099***
						(0.015)
Interaction effect ($\ln m . \ln GTE$)						-0.137*
						(0.069)
Log age-dependency ratio (ADR)	-0.323	-0.350	-0.368***	-0.601***	-0.360***	-0.021
	(0.306)	(0.234)	(0.106)	(0.112)	(0.098)	(0.079)
Interaction effect (ln ADR. $\ln \tau$)				-0.212***		
				(0.038)		
Constant	-0.043	0.047	2.182***	2.801***	1.977***	1.748***
	(1.300)	(1.134)	(0.504)	(0.453)	(0.501)	(0.522)
R-squared	0.677	0.677	0.537	0.616	0.582	0.808

Note: Robust standard errors in parentheses; n=102. *** p<0.01, ** p<0.05, * p<0.1