Preliminary draft; comments on this paper are welcome but the paper is not for circulation

Social Protection and Economic Development: Are the Poorest Being Lifted-Up or Left-Behind?[1]

Martin Ravallion
Department of Economics
Georgetown University,
Washington DC

Dean Jolliffe
Development Data Group
World Bank,
Washington DC

Juan Margitic
Department of Economics
Georgetown University,
Washington DC

Abstract: Standard measures of poverty may reveal nothing about whether the poorest of the poor are being lifted-up or left-behind, yet this is a widespread concern among policy makers and citizens. The paper assesses whether social spending has reached the poorest and hence lifted the floor, and what role economic development plays. Across countries, we find higher floors when mean income is higher. The bulk of this income effect is direct rather than via higher social spending, although social spending generally lifts the floor and is more generous in richer countries. An inequitable growth process led to a sinking floor in the US since around 1990, though stimulus sending on food stamps helped stabilize the floor in the wake of the 2008 crisis. Food stamps lifted the US floor more than one would expect from a uniform (untargeted) allocation of mean spending, but the opposite holds for developing countries, most of which do not lift the floor through social spending any more than the (often-meager) level of mean spending. There are signs of complementarity between social spending and economic growth in determining the gains to the poorest from social spending.

Keywords: Poverty; inequality; floor; economic growth; safety-nets; food stamps

JEL: I32, I38, O15

---

[1] These are the views of the authors and need not reflect those of their employers including the World Bank. Beyond the salary of one of the authors, no financial support for this research was received from the World Bank, or any other organization.

1. Introduction

Development 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] Casual observations of specific antipoverty policies in practice have also looked to the poorest. 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), apparently left behind in the country's rapid economic development. All this 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 principles of justice proposed by Rawls (1971).

But what does it mean to be the "poorest," or to be "left behind"? These expressions can be interpreted as referring to the lowest level of material living and whether it is rising or not. That lower bound can be called the "floor." This should not be confused with the idea of a "biological floor." Human physiology makes it 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 floor. The question is whether, and to what extent, that happens in reality.

Though neglected in the poverty measurement literature, the idea of lifting the floor above the mere biological minimum for survival has long played a prominent role in social policy discussions. Direct interventions have been used against poverty in rich countries and are becoming popular in poorer countries. The policies concerned are of various types and come

[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 United Nations' (2011) *Millennium Goals Report*.) 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." (This was on the occasion of the International Day for the Eradication of Poverty in 2014.) The Vatican's representative to the United Nations also claimed in 2015 that the poorest of the world are being left behind (James, 2015).

under various labels, including "anti-poverty programmes," "targeted interventions," "social safety nets," "social assistance," "social insurance" and "social protection." We shall call them all social protection (SP) and call public spending on SP, "social spending." SP policies are well established policies in rich countries but have not been common in poor ones. That is changing rapidly. SP coverage in the developing worldh has expanded rapidly over the last 20 years, with one or more programs now found in most countries (Ravallion, 2017). In terms of population coverage, the two largest programs are clearly 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 for survival. However, there are continuing concerns that such social protection efforts are not adequately reaching the poorest.[3] Various reasons are given including a lack of political will for the social policies needed, weak adminstrative capacity for policy implementation, ignorance of their rights among poor people, and social stigma associated with targeted antipoverty programs.

The main approach to assessing the distributional impacts of social policies has been to compare the number of poor (or some other income group) before and after the policy intervention, as estimated by comparing the observed gross income distributions with distributions obtained by subtracting reported transfer receipts for each household.[4] This might entail measuring the proportion of the population living below a poverty line or a measure giving higher weight to poorer people. This can be called the "counting approach." In the specific context of assessing poverty impacts, there is a large literature on this approach to poverty measurement, in which various axioms have been proposed. The measures of poverty before and after intervention are obtained by counting poor people, though possibly with higher weight to the poorest. The precise measure of poverty is usually a population-weighted average of individual measures of poverty across the population (counting the non-poor as having zero poverty); Atkinson (1987) characterizes the class of additive measures in the literature and an

[3] With regard to the two examples 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).
[4] This has been the main approach in the literature on benefit-incidence analysis (Kakwani, 1986; van de Walle, 1998). Recent examples include Lindert et al. (2006), Martinez-Vazquez (2008) and the country studies summarized in Lustig et al. (2014). There is also a literature on the effects of social spending on health and education outcomes for which a different approach is required since one cannot simply net out the gains from social spending; examples include Anand and Ravallion (1993), Bidani and Ravallion (1997) and Haile and Nino-Zarazua (2017).

example is the Foster-Greer-Thorbecke (FGT) (1984) measure.

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.[5] 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 social spending roughly halved the average poverty gap index (the agrgegate gap below the poverty line normalized by the line). The counting approach suggests that, as a rule, economic development and social spending 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. For the poorest to not be left behind there must be an increase in the floor. That is what we need to measure, side-by-side with the counting approach. To illustrate the inadequacy of prevailing approaches, Figure 1 shows two pairs of cumulative distribution functions with and without a social protection policy (or before and after economic growth). 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."

From the perspectives of both economic development and social policy there is an interest in knowing more about the level of the floor and how it is evolving. This is not easy, however. There are limitations in 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 (including those in institutions such as worker dormitaries or prisons) are under-represented or even excluded, and these people could be among the poorest stratum. This is a greater concern in advanced countries than in the developing world, but the problem still exists in developing countries; for example, recent rural migrants in cities living in dormitories or slums could well be under-represented. With a sound sampling design and large enough sample we can be reasonably confident about an estimate of the overall mean for those who live in dwellings, but it is far from clear how reliably we could estimate the floor. Ideally one would prefer something like the lower bound of the "time-mean" of consumption or income, measured accurately over a longer period than typical survey data. If we were to know the true

[5] For an overview of the evidence see Ravallion (2016, Chapter 8).

consumption observed over a long enough period in panel data and a large-enough sample we could reliably estimate the floor directly as a time-mean. But that is not the data normally available. And we must recognize the existence of measurement errors in the cross-sectional survey data available for most countries. There are also likely to be transient effects in those data, whereby observed incomes (or consumption expenditures) in a survey fall temporarily below the floor (such as due to illness), but recover later. 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 for observed low incomes is clearly necessary.

Here we follow the approach to measuring the floor in Ravallion (2016b). The essential idea is simple: to estimate the floor by taking a weighted mean of the observed consumptions or incomes (depending on the data) within some stratum that is agreed to be poor, with highest weight on the lowest observed value, and declining weight thereafter within this stratum. However, Ravallion did not study national values of the floor or the role played by 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.

The following section outlines the approach to measuring the floor, while Section 3 discusses some theoretical issues that arise when one introduces social spending, noting that this has both a direct effect on the floor and as a channel linking economic development to the floor. We then turn to the evidence. First we look at the cross-country evidence on the level of the floor across countries, how much it responds to social spending, and the differences between countries in the efficacy of social spending in lifting the floor (Section 4). Then we turn to what has been happening over time in the US where we focus on the Supplemental Nutrition Assistance Program (SNAP)—known as “food stamps” (Section 5), including SNAP’s role in preventing the floor from falling in the aftermath of the financial crisis of 2008. Section 6 concludes.

2. Measuring the floor

We need an estimator for the floor that can be implemented with cross-sectional data, while recognizing that the lowest observed income in such a survey is unlikely to be a reliable indicator. Without loss of generality, we can 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 seems reasonable to assume that this 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 worst off. This assumption can, of course, be questioned. As already noted, the sample of households may have underrepresented (or even excluded) the poorest when they do not in fact live in households. But the assumption appears reasonable when applied to the data in hand. It is also seems reasonable to assume that the probability of being the poorest declines as the person's observed measure of income rises. And beyond some point it would be reasonable to say that there is no chance of being the true floor.

We have an observed distribution of income, y_i, i=1,…n, and a corresponding unobserved distribution y_i^* after eliminating the transient effects and measurement errors. Let the floor of the y_i^* distribution be denoted $y_{min}^* = \min(y_i^*, i = 1, \dots n)$. 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. We can write this as:

$$E(y_{min}^*|y) = \sum_{i=1}^{n} \phi(y_i)y_i \qquad (1)$$

Here the probability that person i, with the observed y_i, is in fact the worst off person is $\phi(y_i)$. This attains its maximum value for $y_{min} = \min(y_i \ , i = 1, \dots n)$ and then falls monotonically as y_i rises, until it reaches zero at some threshold level z, above which there is no chance of someone with that income being the poorest. (In applications, it is natural to use prevailing poverty lines for z, but this is not essential, so we call z a threshold.) The specific functional form satisfying these assumptions proposed by Ravallion (2016b) is:

$$\phi(y_i) = k(1 - y_i / z)^{\alpha} \text{ for } y_i \le z \qquad (2)$$

$$= 0 \text{ for } y_i > z$$

Here there are three parameters, k, α and z, all positive constants. The k parameter assures that the probabilities add up to unity, which requires that $k = 1/(nP_\alpha)$ where P_α is the FGT measure:

$$P_\alpha = \frac{1}{n}\sum_{y_i \le z}(1 - y_i/z)^{\alpha} \qquad (3)$$

Note that the interpretation of α is different to that of the FGT index. Here α determines how fast the chance of being the poorest person falls as y increases rather than the degree of aversion to inequality among the poor, as in the FGT index.

We can then derive the following formula for the expected value of the floor (Ravallion, 2016b):

$$E(y_{min}^{*}|y) = z(1 - P_{\alpha+1}/P_{\alpha}) \tag{4}$$

For example, if we assume that the probability of being the worst off person falls linearly with *y* up to *z*, i.e., that $\alpha = 1$, then the expected value of the floor is $z(1\text{-}SPG/PG)$, where *SPG* is the squared poverty gap index ($\alpha = 2$) and *PG* is the poverty gap index ($\alpha = 1$). Note that $\alpha = 0$ is ruled out by our assumption that the probability falls as *y* increases among those with $y_i \leq z$.[6]

3. Theoretical considerations

On introducing social spending, we now distinguish the pre-transfer floor (y_{min}^{*pre}) from its post-transfer value (y_{min}^{*post}). We are interested in how y_{min}^{*post} varies with both social spending per capita, denoted τ, and overall economic development as measured by the mean *m* of the observed distribution (*y*). We can think of y_{min}^{*pre} as being determined by how the overall mean, *m*, is shared within an economy, while the gain in the floor due to social spending $y_{min}^{*post} - y_{min}^{*pre}$ is determined by how social spending is shared. This points to a separable structure of the form:

$$y_{min}^{*post} = \vartheta(m) + \varphi(\tau) \tag{5}$$

Here $\vartheta(.)$ and $\varphi(.)$ are the sharing functions determining the pre-transfer floor and the gains from social spending respectively. A special case is when the poorest receive the mean social spending. One way this could happen under a universal basic income (UBI) for which everyone receives τ. Then we have what we term the "UBI null hypothesis" that $\varphi(\tau) = \tau$, i.e.,

$$y_{min}^{*post} = y_{min}^{*pre} + \tau \tag{6}$$

Our empirical analysis will follow standard practice in benefit-incidence analysis, of estimating the pre-transfer distribution by subtracting transfers received. This ignores behavioral responses such as through saving or labor supply. In defense of this assumption, it can 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 can provide a partial test of this assumption, which can be thought of as a consistency check. The reasoning is as follow: if there are behavioral responses by the poorest then y_{min}^{*post} less the transfers received by

[6] If one uses $\alpha = 0$ then every consumption below *z* is deemed equally likely to be the lowest, so $z(1\text{-}PG/H)$ is the mean for the poor, where *H* is the headcount index.

the poorest τ_{min} will equal the true y_{min}^{*pre} plus a behavioral effect, the bias b. If $b > 0$ then we can expect that its value tends to rise with the overall level of social spending. Specifically, the test assumes that: (i) true value of y_{min}^{*pre} is a function of the mean, *m*, and (ii) the bias b is a function of mean spending:

$$y_{min}^{*post} - \tau_{min} = y_{min}^{*pre}(m) + b(\tau) \qquad (7)$$

The test is then to see if there is a partial correlation between the estimated pre-transfer floor and mean spending at a given value of mean income. Intuitively, when it is correctly measured, y_{min}^{*pre} should not vary with the level of social spending at a given mean income.

Separability must also be considered a strong assumption in this context. For example, a higher mean income may well come with administrative capabilities (including information systems) that allow governments to better reach the poorest. To see how, suppose that economic development comes with structural changes such that a rising share of national income is derived from formal-sector activities that are 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 comes with formalization, generating greater administrative capability including for effective social spending. Then it can be expected that economic development allows higher social spending and supports a greater capacity to make that spending effective in reaching the poorest. To give another example, lack of knowledge about how to access public programs has often been identified as a factor weakening the coverage of social protection policies for poor people.[7] 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 social spending will tend to rise with mean income when comparing different countries.

So instead of (5) we write this relationship of interest in the more general form:

$$y_{min}^{*post} = f(\tau, m) \qquad (8)$$

Here f is some (smooth) function. So the pre-transfer floor is $y_{min}^{*pre} = \vartheta(m) = f(0, m)$. In the empirical work we test separability, to see if there is a degree of complementarity between economic development and social spending; specifically, does a higher mean income increase

[7] See, for example, Ravallion et al. (2015) in the context of a large workfare program in India and Daponte et al. (1999) on the context of food stamps in the US.

the marginal gains to the poorest from higher social spending? Such 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 social spending on the floor stem from differences in the overall level of social spending versus differences in transfer efficiency. For this purpose we measure what we term Floor Transfer Efficiency (FTE), defined as:

$$FTE = (y_{min}^{*post} - y_{min}^{*pre})/\tau \tag{9}$$

A higher FTE does not of course imply a larger impact on the floor, for which $y_{min}^{*post} - y_{min}^{*pre}$ is the relevant indicator. 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 share of total transfers received by the poor, which is the reduction in the aggregate poverty gap per $ spent.[8]

Economic development and the floor: We can identify two channels in how economic development can impact the floor. The first is direct, in that it holds at any given level of social spending. This channel arises through the distribution of the income gains associated with economic growth. Intuitively, the more "pro-poor" the growth process—the more it augments demand for relatively unskilled labor—the stronger will be this direct channel. However, being "pro-poor" is not the same thing as reaching the poorest, as discussed in the Introduction and Section 2; poverty measures can fall yet the poorest are left behind. Indeed, we may see a sinking floor with certain growth processes. Suppose that growth is generated by greater trade openness and technological change both of which put downward pressure on unskilled wages and hence the floor. Then we could see the floor fall with a higher mean.

The second channel is indirect, via higher social spending. As has long been recognized in discussions of poverty and growth, a potentially important channel by which economic growth can reduce poverty is via higher social spending.[9] But is this channel important in practice, and does it embrace the poorest? The growth may be heavily concentrated among an elite who use their economic power to further reinforce their positions by supporting political opposition to redistributive tax and spending policies, with implications for the poorest as well as others.

[8] This measure 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.
[9] See, for example, the discussion in Anand and Ravallion (1993). For an overview of these arguments, see Ravallion (2016a, Chapter 8).

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 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. And we will assess the relative importance of the direct and indirect channels.

A simple theoretical model of the political economy of the indirect channel provides some insights regarding this indirect channel. (This is not the only model one can write down, but it will suffice for this purpose.) Let us assume that the aggregate level of social spending is chosen by the median voter who faces a trade-off between the current tax burden to finance social spending and the expected future protection, given that there is a chance of falling to the floor in the event of a shock.[10]

What then 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 nonetheless sway electors. Against this view, the observed distribution also includes measurement errors that would be less likely to matter to electoral outcomes. Here we will assume that the relevant median, denoted y^{med}, is that of the observed distribution. The model can be readily modified to allow the alternative assumption that it is the median of the y_i^* distribution that matters. Of course, our empirical work has no choice but to use the observed medians.

A uniform tax τ is levied to finance social policies, which depletes the current net income of the median voter. There is a positive probability that the median voter will fall to the floor in the future, and the voter takes account of the effects of higher social spending on the level of that floor. The utility function is $u(y)$ which is strictly increasing and concave. The median voter choses τ to maximize:

$$u(y^{med} - \tau) + \rho u(y_{min}^{*post}) \quad (10)$$

[10] The median voter model is not the only way one can think about this; an option is to postulate a governmental social welfare function.

Where ρ is the discount factor, which can be thought of as the probability of falling to the floor in the future. We require that $0 < \rho < 1$. We allow the possibility that $f_{\tau\tau} > 0$, but that this is bounded above such that:

$$\frac{f_{\tau\tau}}{f_\tau^2} < \frac{-u_{yy}\left(y_{min}^{*post}\right)}{u_y\left(y_{min}^{*post}\right)} \tag{11}$$

When combined with our assumption that *u*(*y*) is strictly concave for all *y*, the condition in (11) assures that the second-order condition for a unique optimal level of social spending is satisfied.

The median voter's optimal level of social spending solves the first-order condition:

$$u_y(y^{med} - \tau) = \rho u_y(f(\tau, m))f_\tau(\tau, m) \tag{12}$$

We can write the solution as:

$$\tau = \tau(y^{med}, m) \tag{13}$$

with first derivatives:

$$\tau_{y^{med}} = \frac{u_{yy}(y^{med} - \tau)}{u_{yy}(y^{med} - \tau) + \rho[u_y\left(y_{min}^{*post}\right)f_{\tau\tau} + u_{yy}\left(y_{min}^{*post}\right)f_\tau^2]} > 0 \tag{14.1}$$

$$\tau_m = \frac{-\rho[u_y\left(y_{min}^{*post}\right)f_{\tau m} + u_{yy}\left(y_{min}^{*post}\right)f_\tau f_m]}{u_{yy}(y^{med} - \tau) + \rho[u_y\left(y_{min}^{*post}\right)f_{\tau\tau} + u_{yy}\left(y_{min}^{*post}\right)f_\tau^2]} \tag{14.2}$$

While $\tau_{y^{med}} > 0$ (given (10) and $u_{yy} < 0$), the sign of τ_m is ambiguous. A key issue is the degree of complementarity between social spending (higher τ) and economic development (higher *m*) in raising the floor, as indicated by the cross-partial derivative $f_{\tau m}$. This 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 social 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:

$$\frac{f_{\tau m}}{f_\tau f_m} > \frac{-u_{yy}\left(y_{min}^{*post}\right)}{u_y\left(y_{min}^{*post}\right)} \tag{15}$$

If this condition holds then we will say that there is strong complementarity between economic development and social spending in how they influence the level of the floor. It is evident from (14.2) that strong complementarity implies that $\tau_m > 0$.

When we consider the bivariate relationship between social spending and economic development we need to bring in the effect of a higher mean on the median. Intuitively, the higher m is relative to y^{med} the higher is inequality. (Indeed, we can think of m/y^{med} as an indicator of inequality.) The total effect of economic development on social spending is:

$$\frac{d\tau}{dm} = \tau_m + \tau_{y^{med}} \frac{dy^{med}}{dm} \tag{16}$$

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

Case 1: Equitable growth has both a direct and indirect gain to the poorest. In this case, the growth process lifts the median, $\frac{dy^{med}}{dm} \geq 0$. Suppose also that there is strong complementarity. Then the effect on the floor is:

$$\frac{dy_{min}^{*post}}{dm} = f_m + f_\tau \frac{d\tau}{dm} > 0 \tag{17}$$

If the function $f(\tau, m)$ only exhibits weak complementarity (or substitutability) between economic development and social spending 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:

$$\rho[u_y(y_{min}^{*post})f_{\tau m} + u_{yy}(y_{min}^{*post})f_\tau f_m] - u_{yy}(y^{med} - \tau)\frac{dy^{med}}{dm} > 0 \tag{18}$$

Case 2: Inequitable growth leaves the poorest behind. If there is only weak complementarity and economic development is sufficiently inequality increasing then it is possible to find that neither social spending nor the level of the floor respond positively to a higher mean income. Suppose that there is only weak complementarity, i.e., $\frac{f_{\tau m}}{f_\tau f_m} \leq \frac{-u_{yy}(y_{min}^{*post})}{u_y(y_{min}^{*post})}$, and that economic growth (a high m) does not benefit the median voter ($\frac{dy^{med}}{dm} = 0$). Then social spending falls with economic development, and it is also possible that the floor does not rise, and even falls if $f_m + f_\tau \tau_m < 0$.

4. Cross-country evidence

We now take the ideas of the last two sections to a cross-country empirical setting. For this section we mainly rely on the data available in the World Bank's "The Atlas of Social Protection" (ASPIRE) as accessed mid-2017.[11] This dataset contains a set of indicators created to

[11] ASPIRE is essentially a cross-country compilation of the outputs from a software program produced by the World Bank, *ADePT Social Protection*. Tesliuc and Leite (2010) provide a user manual.

help assess the performance of governmental social protection systems (social assistance, social insurance and active labor market programs); the bulk of this (81% on average) is social insurance.[12] ASPIRE draws on 262 household level surveys in 122 countries, mostly in the developing world, from 1998 to 2014.[13] This database offers a variety of useful indicators that allow us to account for the range of social safety nets, their type and coverage. We also include GDP per capita from the World Bank's *World Development Indicators* database. All required currency conversions are done at 2005 PPP.

The ASPIRE data set includes the poverty measures *PG* and *SPG* using poverty lines *z* fixed at the quantile of *H*=20%. We use these two measures to construct our estimates of floor for $\alpha = 1$.[14] In the ASPIRE dataset, the computations for *SPG* and *PG* pre-transfer are done assuming no behavioral responses. The ASPIRE dataset provides total social protection transfers per capita, which we use as our measure of τ.

Summary statistics and graphs: Table 1 provides summary statistics. The (un-weighted) mean floor pre-transfers is \$1.21 a day, though varying widely, from \$0.03 to \$4.82. There is undoubtedly measurement error; it is very hard to believe that anyone lives at \$0.03 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.

We can observe from Table 1 that social spending lifted the floor by \$0.52 a day on average, 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 (a standard error of 0.059). Thus, we can reject the "UBI null hypothesis" implied by (6). We also observe in Table 1 that social spending reduced the headcount index by about 7% points (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%.

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

[13] 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 social 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 is almost certainly a data error.

[14] The alternative of fixing the poverty line at (say) the World Bank's international line was rejected as it yields very small subsamples for estimating the floor in many countries.

We find that 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 social spending and differing transfer efficiencies. Transfer efficiency in reaching the poorest varies greatly. Figure 4 gives the empirical density function for FTE.[15] (Recall that this is the ratio of the gain in the floor due to social spending to mean spending.) We see that very few countries attain a FTE of unity or more. For the bulk of countries (87% of the sample), the gain to the poorest is less than mean social spending.

In addition to allowing us to measure the floor before and after transfers, and hence FTE, the ASPIRE data set allows us to measure the efficacy of social spending in reaching the poorest 20%. Recall that we have two measures, FTE and GTE. The two are correlated (r=0.505), but possibly not as much as one might have thought; some countries are better than others at reaching the poorest people given their efficacy in reaching the poorest 20%. The value of GTE is positively correlated with spending per capita (r=0.656), but this 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 we find with average social spending for both FTE and GTE (it is easier to see if one logs social 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 the mean (r=-0.476).

However, the bulk of the variance in the effectiveness of social spending in reaching the poorest (as evident in Figure 3) is due to the variance in levels of social spending rather than the efficiency of that spending 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 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.)

It will be recalled that our theoretical model in Section 3 suggests that a key issue is how much the median rises with the mean. We find a positive relationship. Indeed, the OLS elasticity of the median to the mean—the regression coefficient of log median on log mean—is not significantly different from unity; the regression coefficient is 1.012 with a robust standard error

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

of 0.017. Under the assumptions of our median-voter model in Section 3, we then expect mean transfers to rise with the overall mean. This pattern is evident in Figure 6 showing how log mean transfers vary with the log mean.

We see in Figure 7 that countries with a higher median tend to have higher social spending, as predicted by our theoretical model. Given the pattern in Figure 6, it is no surprise that the bivariate relationship is very similar with the mean, though as we will see below, the stronger partial correlation is with the median, once one controls for the mean.

Richer countries tend to have a higher floor. Figure 8 plots the relationship between the (log) post-transfer floor and (log) mean. There is a strong positive relationship across countries between the level of the floor and mean income. In terms of the model in Section 3, this reflects both higher social spending in richer countries, and a direct effect at given spending. We will now separate out these effects.

Partial correlations: To allow for multiple covariates, we next explore these relationships further using regressions. (These are not intended to be given a casual interpretation, but are only a convenient means of testing for partial correlations.)

First, we test for behavioral responses by the poorest. Recall that an implication of our assumption that the pre-transfer floor is the post-transfer value less social spending received by the poorest is that we should not find a correlation between the estimated pre-transfer floor and mean spending (Section 3). While there is a significantly positive (zero-order) correlation between the pre-transfer floor and average social spending (r=0.511), this vanishes when we control for the mean. The partial correlation falls to 0.068; Figure 9 plots 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 elasticity is about 0.8 and it 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. The restriction that social 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.

Next, we reexamine how social spending varies with the mean, but now controlling for the median. Recall that our model of the political economy of social spending implies that the median matters independently of the mean (equation 13), but only the comparative static effect of the median is predicted in sign. The expected positive effect is confirmed by the data; Table 2 gives the relevant regression. There is a strong positive coefficient on the (log) median, while the (log) mean has a negative effect. This is not consistent with complementarity (Section 3). However, if we use GDP per capita instead of the survey mean, the coefficient is positive.

Next, we look at the post-transfer floor. Table 3 gives the corresponding regressions for both the post-transfer floor and the gain in the floor attributed to social spending, i.e., our estimate of $\ln(y_{min}^{*post}/y_{min}^{*pre})$. The post-transfer floor has an elasticity of 0.9 to mean income (Column 4). This is not significantly different from unity, implying that the floor does not fall relative to the mean as economies develop. Thus we see that social spending is able to negate the tendency for the pre-transfer floor to fall as a share of the mean with economic development. This comes from the indirect effect of a higher mean via social spending, although the bulk (74%=0.686/0.923) of the effect is direct. We also see that higher aggregate transfers contribute to a larger impact of transfers on the floor (Columns 5 and 6). Countries with a higher mean income tend to see somewhat lower impacts of SP transfers on the floor, though this is only significantly different from zero at the 6% level.

We find that there is a positive interaction effect between average transfers and the mean, which helps in raising the impact of social spending on the floor (Table 3, Column 8). There is complementarity between social spending and economic development in how they influence the efficacy of social spending in raising the floor.[16]

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 social 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 (9) of Table 3.

As expected, there is a strong interaction effect between GTE and transfer spending in their effects on the extent to which social spending lifts the floor. There is also a negative interaction effect between the survey mean and transfer efficiency; effectiveness in transferring money to the poorest 20% matters more to reaching the poorest of the poor in poor countries.

[16] Note that $\frac{\partial^2 \ln y_{min}^*}{\partial \ln m \, \partial \ln \tau} > 0$ implies $\frac{\partial^2 y_{min}^*}{\partial m \, \partial \tau} > 0$.

When we evaluate the total effects at the mean points, we find a significant positive effect of social spending and GTE on the extent to which those transfers succeed in raising the floor (Table 3, lower panel). Once we control for the level of transfers and transfer efficiency we do not find evidence that higher average incomes come with a greater impact of transfers on the floor.

Note that Table 3 uses the survey means. The qualitative results are similar if one uses GDP per capita instead. The most notable difference is that the mean income elasticity of the floor is lower when one uses GDP while the elasticity of the post-transfer floor to social sending is somewhat higher. Controlling for (log) social spending, the elasticity of the pre-transfer floor w.r.t GDP per capita drops to 0.453 (s.e.=0.050) while that for the post-transfer floor w.r.t. GDP drops to 0.615 (0.048), and to 0.256 (s.e.=0.077) when one controls for social spending. The social-spending elasticity controlling for (log) GDP per capita rises to 0.204 (0.039). A positive interaction effect between social spending and mean income is also evident when one switches to GDP; this holds for both the (log) post-transfer floor and the proportionate gain in the floor.

5. Evidence for the US

While the patterns found in the cross-country data are at least suggestive of how the floor may evolve over time, it is of obvious interest to look at actual time-series evidence.

One reason for choosing the US for this purpose is the availability of annual time series on the relevant variables. Another reason is that (as is well known) there has been a strong tendency of rising inequality, and limited progress against poverty pre-transfers. We focus on the largest direct intervention against poverty in the US, the food stamps program (SNAP). SNAP is a Federal program (administered by the Department of Agriculture) that helps poor families purchase food. In (fiscal) 2016, SNAP covered about 44 million Americans (14% of the population) at a cost of $71 billion, representing $125 per person per month for food stamp beneficiaries. The program is targeted to poor families (living below 130% of the official poverty line), and aims to provide larger benefits to poorer families.[17] There are concerns that the program does not reach all those who are eligible.[18]

[17] For further information and analysis on SNAP see Bartfeld et al. (2016). On the benefits to children from poor families see Jolliffe et al. (2005).
[18] See, for example the discussion in The Economist (2011).

These time series data are unlikely to have much power for testing the political economy model of social spending in Section 3, and so we will not try to explain SNAP spending by the time series variation in median income (though this may still be an underlying longer-term property of the data). There are, however, two significant policy changes worth noting. First, a series of reforms in 1996-98 put emphasis on reducing perceived "leakage" to those not considered eligible, including greater use of work requirements, though this can also reduce participation by eligible participants including some of the poorest.[19] Time limits and recertification became stricter; legal immigrants were variously eligible, then ineligible, then eligible again, but growing concerns about status reduced their participation. Able-bodied adults without dependents found it harder to access SNAP.[20] Second, spending on the program surged in the aftermath of the 2008 crisis and the subsequent rise in unemployment. As a part of the American Recovery and Reinvestment Act, SNAP benefits increased by 14% in April 2009.

To assess the effects of SNAP on the floor, we use the micro data from the Annual Social and Economic (ASEC) Supplement to the Current Population Survey (CPS). We use 26 years of CPS-ASEC data from 1989 to 2014, which allows us to estimate poverty measures and SNAP benefit levels from 1988 to 2013. The CPS is administered monthly by the Census Bureau for the Bureau of Labor Statistics and collects data from a nationally representative sample of households on employment, unemployment, earnings, occupation, and hours of work. Respondents provide information on several different sources of income, including noncash income sources such as SNAP. The CPS is also the data source for official US poverty measures, and we use the same measure of income (pre-tax money income excluding capital gains, noncash benefits and tax credits) as is used in the official measures of poverty. To the reported value of income, we add the face value of the food stamp (or in later years of our data, the reported credit value placed on the electronic benefit transfer card) to obtain post-SNAP incomes.

Unlike the cross-country analysis for which we assumed that the floor is within the poorest 20%, for the US we assume that the poorest is below the official US poverty line (fixed

[19] The first reform was the Personal Responsibility and Work Opportunities Reconciliation Act of 1996 (with most provisions only effective from mid-1997), followed by the Balanced Budget Act of 1997 and the Agricultural Research, Education and Extension Act of 1998 (USDA, 2017). For further discussion of these reforms to SNAP and their implications see Currie and Grogger (2001).

[20] Electronic Benefit Transfer (EBT) was also introduced around the turn of the century (earlier in some states), such that SNAP recipients paid for food using a "debit card." (If the pin code is verified and the account balance is adequate then payment is accepted.) It is unclear on a priori grounds whether this would help the poorest, although the bulk of the decline in FTE appears to have preceded EBT.

in real terms). But this is not such a big a difference given that the US official poverty rate has changed very little over the period; it was 13% in 1988 and 14% in 2013. Fixing the poverty rate (instead of the poverty line) at say 15% would probably give very similar results.

Summary statistics: Table 4 provides the key data by year and summary statistics. The mean post-SNAP floor is about 36% of the official poverty line, which was $24,036 a year in 2015 for a family of four (two adults and two children) or $16.50 per person per day (rounding up slightly). So the mean floor in that year's prices is $5.89 a day, while the pre-transfer value is $5.40. Unlike the cross-country data set, the mean gain to the poorest now exceeds mean spending on food stamps; the difference between the mean gain for the poorest and mean spending of $0.23 a day is statistically significant (s.e.=0.037). So (again) the null implied by (6) is rejected, but this time the gain to the poorest exceeds mean spending.

Similarly to Figure 6 (using the cross-country data), we find that the median increased with the overall mean. However, the ratio of the median to the mean has been falling over time in the US, from around 0.8 at the beginning of the period to 0.7 at the end (Table 4). While the median rises with the mean, as in the cross-country data set and consistently with our model in Section 3, unlike the cross-country data, it does so with an elasticity less than unity.[21] This decline in the median relative to the mean as the latter increases is another aspect of the inequitable growth process of the US. The Gini index of family incomes rose from 0.40 to 0.48.

We see from Table 4 that food stamps increased the floor by 0.03 on average (as a proportion of the official poverty line). We also see from Table 4 that SNAP benefits per capita rose with either a higher mean or median. Furthermore, it does so with an elasticity of about unity; the regression coefficient of log food stamp spending on log mean is 1.009 (s.e.=0.323) while it is 1.167 (0.359) for the log median. Figure 10 plots the series on food stamp spending over the period. Per capita spending increasing by a factor of about four between 2000 and 2013, with the bulk of this being due to the sharp rise in the aftermath of the financial crisis, with spending peaking in 2011. This "SNAP stimulus" was part of the American Recovery and Reinvestment Act (ARRA) of 2009.[22]

[21] The OLS regression coefficient of the log median on the log mean is 0.878 with a standard error of 0.013. One can reject the null that the coefficient is unity with t=9.55.

[22] SNAP benefits rose by 13.6% in April 2009. For example, the maximum SNAP benefit for a household of three rose from $463 to $526 per month. Under ARRA, states could suspend time limits for unemployed able-bodied adults. Subsequent legislation imposed an expiration date of November 2013 for the 13.6% SNAP benefit increase.

Figure 11 plots our estimated poverty measures before and after SNAP while Figure 12 plots our estimates of the floor, for both $\alpha = 1$ and $\alpha = 2$. (Naturally, the level of the floor is lower for $\alpha = 2$, whereby the probability of a given income being the floor declines as a quadratic rather than linear. But the trajectories over time are similar, and r=0.998.) We find that food stamps lifted the floor, and by an amount that exceeds overall SNAP spending per capita.

Note that the standard poverty measures are not highly correlated with the floor. For example, the period 1993-2000 saw declining poverty measures but also a sinking floor. And the sharp rise in poverty measures in the crisis period (2008-10) came with a relatively stable floor post-SNAP. The proportionate changes over time (first differences in logs) are roughly orthogonal for the (post-SNAP) headcount index (r=0.056) and only not highly correlated for the two poverty gap indices (r=-0.268 for changes in PG and r=-0.510 for SPG). Tracking standard poverty measures alone is not very informative about what is happening to the floor.

We find that there has been a substantial decline over time in FTE—the ratio of the gain in the floor that we attribute to SNAP to spending per capita on the program (Figure 13). The program is reaching the poorest, but transfer efficiency in raising the floor has declined appreciably over time; the bulk of that decline was in the first 10 years. In the late 1980s, FTE was high, at around 5, meaning that the gain in the floor was five times mean spending, but by the last five years of the series it had fallen to about unity. Food stamps used to be much more effective in reaching the poorest. The decline in transfer efficiency in raising the floor largely preceded the large expansion in SNAP spending under ARRA in the latter sub-period (including the crisis) but the latter expansion clearly did not come with better performance in reaching the poorest.

Over the period as a whole, transfer efficiency is negatively correlated with the level of SNAP spending (r=-0.656). However, there are two distinct periods, before and after 1998, with very different relationships, as can be seen clearly in Figure 14. From 1998 the scheme became less transfer efficient in lifting the floor at a given level of spending, and FTE became less responsive to SNAP spending. The "SNAP stimulus" entailed a predictably lower level of FTE, based on the post-1997 model (Figure 14). Although attribution cannot be proven conclusively, the pattern in Figure 14 is at least consistent with the view that the program's effectiveness in reaching the poorest per $ of outlay declined with the policy changes.

Regressions: The following discussion will focus on $\alpha = 1$ in keeping with the cross-country analysis. Given that the US data are a time series we include the lagged dependent variable and we provide heteroscedasticity and autocorrelation consistent (HAC) standard errors.[23] Table 5 provides the regressions corresponding to Table 3 for the cross-country data set. In contrast to the latter, food stamp spending in the US does not exhibit a positive mean income effect; regressing log food stamp spending on its own lagged value and the log of the mean, the coefficient on the latter is 0.062 with a HAC standard error of 0.122. So in the case of food stamps in the US there is only the potential for a direct effect on the floor.

Consistently with our results for the cross-country data set, we find that higher food stamp spending does not significantly alter the pre-SNAP floor.[24] This is again reassuring about our method. We find that food stamps lifted the floor at given mean income.

A major difference between the cross-country results and those for the US is that economic growth in the US has come with a lower floor. (There is also a small negative effect on the gains from SNAP.) As discussed in Section 3, this is interpreted as the combination of (i) only weak complementarity (or even substitutability) between social protection and economic development with (ii) a distributional effect, whereby growth has come with changes in distribution that have gone against the poorest. When we add the Gini index to the regressions for the pre-SNAP floor (Columns 1 and 2 of Table 5) the coefficient on the mean drops in size and is no longer significantly different from zero.

In Table 6 we test for interaction effects between the mean and SNAP spending in determining the gains to the poorest. (The lagged dependent variable was no longer significant so we dropped it.) Similarly to the cross-country sample, strong interaction effects are evident.

The above results suggest that the expansion of SNAP in the 2000s has at least partly compensated for the downward pressure on the floor coming from the unequal nature of America's growth process. To quantify this effect, we simulate a counterfactual trajectory of the floor, assuming that food stamp spending stayed at its level in 2000 instead of the increase actually observed (Figure 10). Nothing else changes in this counterfactual simulation (including the residuals). We use the specification in Column 2 of Table 6. Figure 15 provides the

---

[23] Without the lagged dependent variable as a regressor, the regressions exhibit significant autocorrelation in the errors, as revealed by (say) the Durbin-Watson statistic. We also tested the robustness of each regression to adding an independent time trend. The key findings were robust to adding a time trend (with one exception noted below).
[24] This also holds when one regresses the first differences of the floor variables on the first difference of log food stamp spending and that for the log mean.

simulation. While (as we saw in Figure 12) the floor stabilized (though with fluctuations) after 2000, without the increase in public spending on SNAP our results imply that it would have fallen. Extra food stamps helped assure that the poorest could at least maintain their (low) living standards.

6. Conclusions

To test whether social spending has lifted the floor we must be able to measure and monitor that floor in a comprehensive and systematic way. It is clearly not enough to look at the evolution of any standard poverty measure, which can fall with or without an increase in the level of living of the poorest. Instead, we need a measure of the floor—the lower bound of the distribution, and we must recognize that the lowest observed income in a cross-sectional survey need not be a good indicator of the true floor of living standards given transient real effects and measurement errors in the data. To address these concerns we have measured the floor as a weighted mean for those deemed to be poor, with highest weight on those observed to be poorest (following Ravallion, 2016b).

To help motivate our empirical analysis, we have outlined a simple model of the determination of the floor, which is taken to depend on both the level of social protection spending and economic development, measured by mean income, while the level of social 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 spending. A key role is played by the extent of complementarity between social spending and development; complementarity exists when higher social spending increases the marginal gains to the poorest from growth. 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 that need not hold, and there is no guarantee that the poorest will see any gain from overall economic development.

We have assembled evidence from both cross-country data for developing countries and time series data for the US. We find that higher social spending helps to lift the floor in both data sets. 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 social spending rather than the transfer efficiency of that spending. However, for the cross-country sample, the gain to the poorest—those living at the floor—from social spending is significantly less than aggregate spending per capita. The opposite holds for food stamps in the US, for which the gain to the poorest is significantly greater than mean spending, though substantially less so over time. Thus, for developing countries, the poorest would do better with a universal basic income anchored to existing social spending, while for the US the opposite is true with regard to food stamps, which are raising the floor more than mean spending would imply. However, the marked decline in the efficacy of food stamps in reaching the poorest since the 1990s suggests that today the gain to them would not be much different from a UBI.

We also 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 with overall growth. The bulk of the efficacy of economic growth in lifting the floor is direct, rather than via higher social spending. Nonetheless, social spending comes close to assuring that the post-transfer floor does not sink relative to the mean when comparing low and high-mean 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 social spending and mean income in regressions for the gains to the poorest from higher social spending. This holds for both the cross-country sample and the US over time. Along with rising social spending, this complementarity plays a positive role in helping to assure that the poorest benefit from growth.

However, in contrast to the cross-country sample, for the US we find that economic growth has come with a sinking floor. We interpret this as the combined effect of relatively weak complementarity with a quite strong distributional effect. The floor in the US has fallen during the 1990s, stabilizing in the 2000s. The expansion of the food stamps program in the wake of the financial crisis was able to prevent a fall in the floor despite the inequality-increasing growth process.

We warn against causal interpretations. Our estimates of the impacts of social spending on the floor can only be interpreted as the true causal impacts under the assumption that there are no behavioral responses by the poorest. We have shown that our estimates of the pre-transfer floors under this assumption are uncorrelated with mean transfers (at a given mean income). This is consistent with our assumption of behavioral neutrality for the poorest, but it is not conclusive evidence. There may still be some bias in our estimates due to behavioral responses by the poorest. We also warn against giving our regressions a casual interpretation; we provide them as only descriptive tools for measuring (partial) correlations.

In thinking about further work on this topic, the greatest concern in any attempt to give our results a causal interpretation is possibly not behavioral responses by the poorest but rather the endogeneity of social spending. Our regression error term includes the (likely) heterogeneity across countries in the impact of social spending on the floor. The regression coefficients we have estimated are interpreted as average impacts. Suppose that countries that are more effective at using aggregate social spending to reach the poorest tend to spend more on transfers. Then our variable for spending on transfers will be positively correlated with the error term, implying that OLS overestimates the true mean impact of that spending on the floor. However, we cannot rule out the possibility of the opposite direction of bias, notably if the countries that are better at using aggregate social spending to reach the poorest tend to spend less on transfers.

Identifying behavioral responses of the poorest and impacts of social spending on the floor remains a challenge, though probably one that is better addressed by a more micro-empirical and program-specific approach using impact evaluation tools. We hope that this study has at least pointed to the desirability and feasibility in such efforts of including a focus on the living standards of the poorest. Doing so will remove an obvious disconnect between how social policy makers view their objectives and how economists assess their progress. Greater interest in measuring the floor will hopefully also come with extra effort in survey design and processing to more accurately measure the living standards of the poorest.

References

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.

Bartfeld, Judi, Craig Gundersen, Timothy M. Smeeding, and James Patrick Ziliak, eds., 2016, *SNAP Matters: How Food Stamps Affect Health and Wellbeing*. Studies in Social Inequality. Stanford, California: Stanford University Press.

Bidani, Benu, and Martin Ravallion, 1997, "Decomposing Social Indicators Using Distributional Data," *Journal of Econometrics* 77(1): 125-140.

Currie, Janet, and Jeffrey Grogger, 2001, "Explaining Recent Declines in Food Stamp Program Participation." *Brookings-Wharton Papers on Urban Affairs*, 203–44.

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.

Dasgupta, Partha, 1993, *An Inquiry into Well-Being and Destitution*, Oxford: Oxford University Press.

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, 2011, "The Struggle to Eat," *Economist Magazine*, July 14.

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

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

Haile, Fiseha, and Miguel Nino-Zarazua, 2017, "Does Social Spending Improve Welfare in Low-Income and Middle-Income Countries?" *Journal of International Development* (Published online).

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

Jolliffe, Dean, Craig Gundersen, Laura Tiehen, and Joshua Winicki, 2005, “Food Stamp Benefits and Child Poverty,” *American Journal of Agricultural Economics* 87(3): 569-581.

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

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.

Ravallion, Martin, 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.

United States Department of Agriculture (USDA), 2017, "A Short History of SNAP," United States Department of Agriculture.

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

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

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

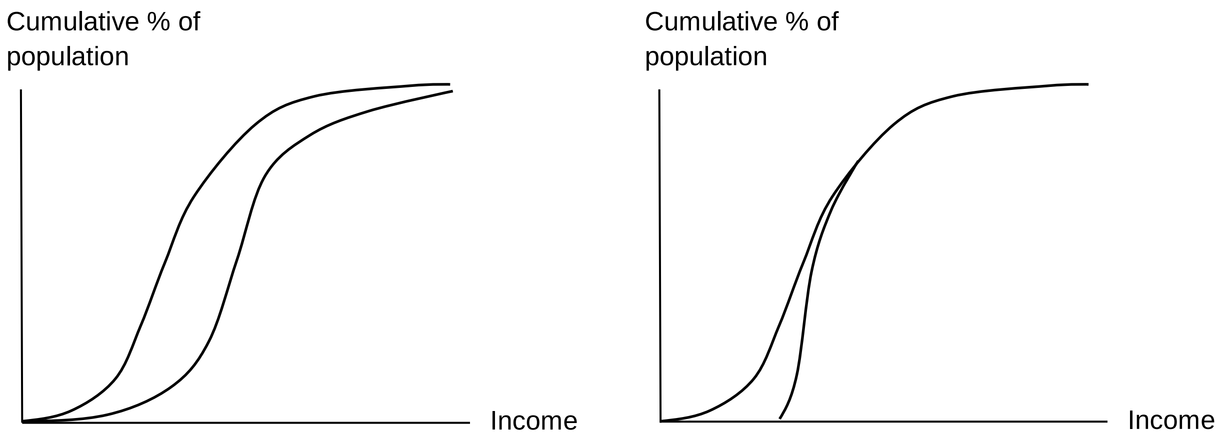

(a) The floor has not risen; poorest are left behind

(b) Floor has risen

Source: Ravallion (2016a).

Figure 2: Kernel density functions for the floor across countries

(a) Linear

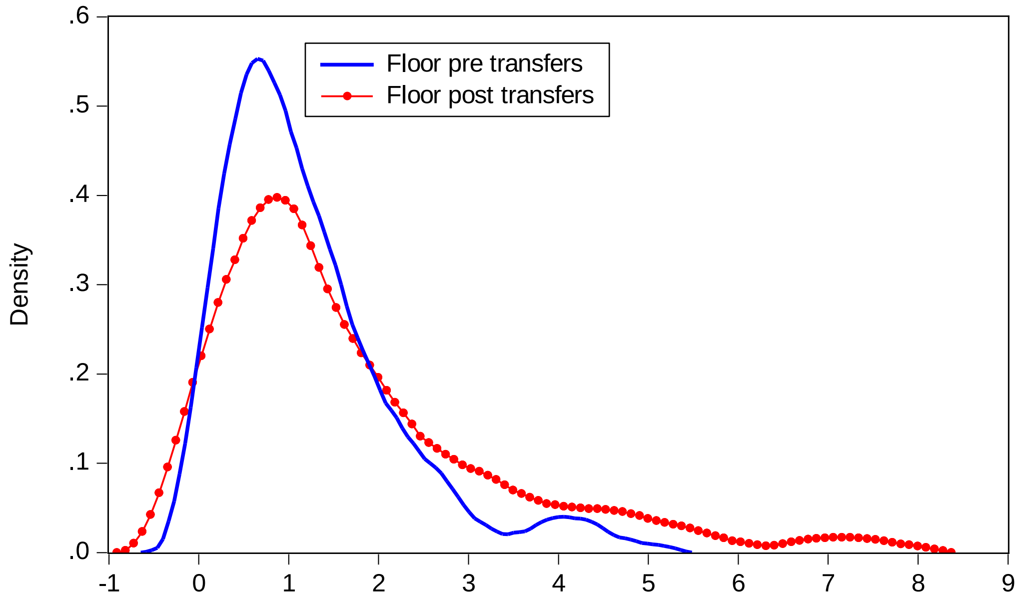

(b) Log transformation

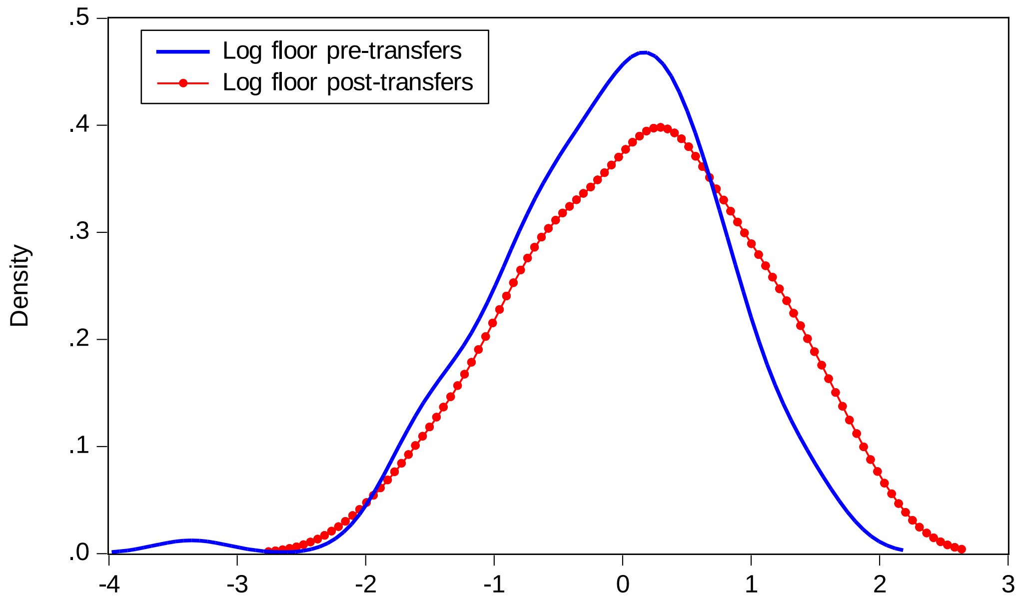

Figure 3: Higher social spending comes with a higher floor

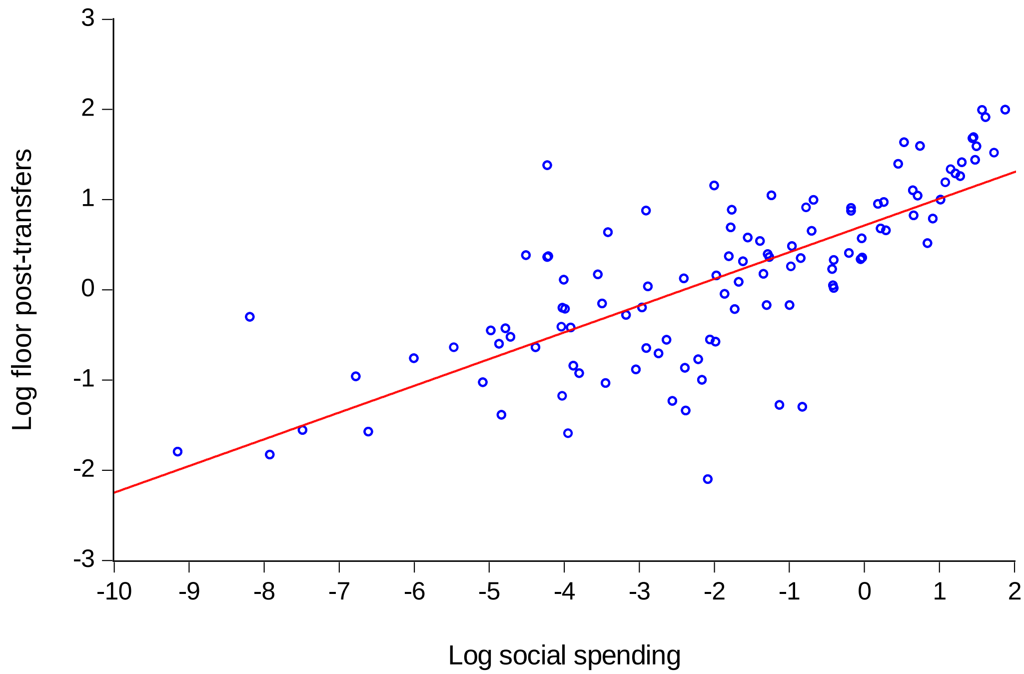

Figure 4: Kernel density functions for floor transfer efficiency

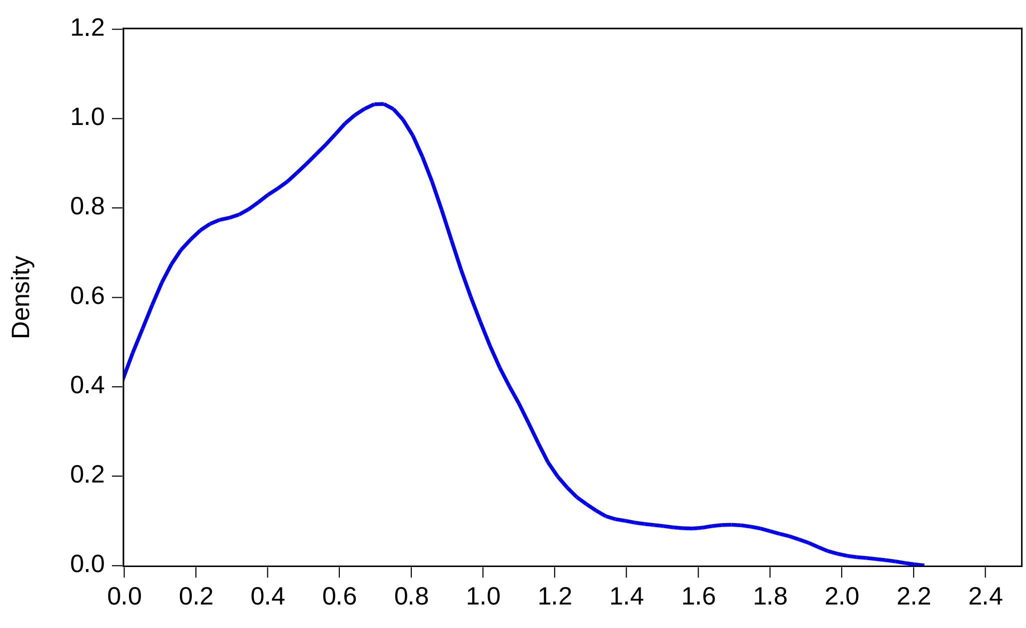

Figure 5: Transfer efficiency plotted against aggregate transfers per capita

(a) Gap transfer efficiency

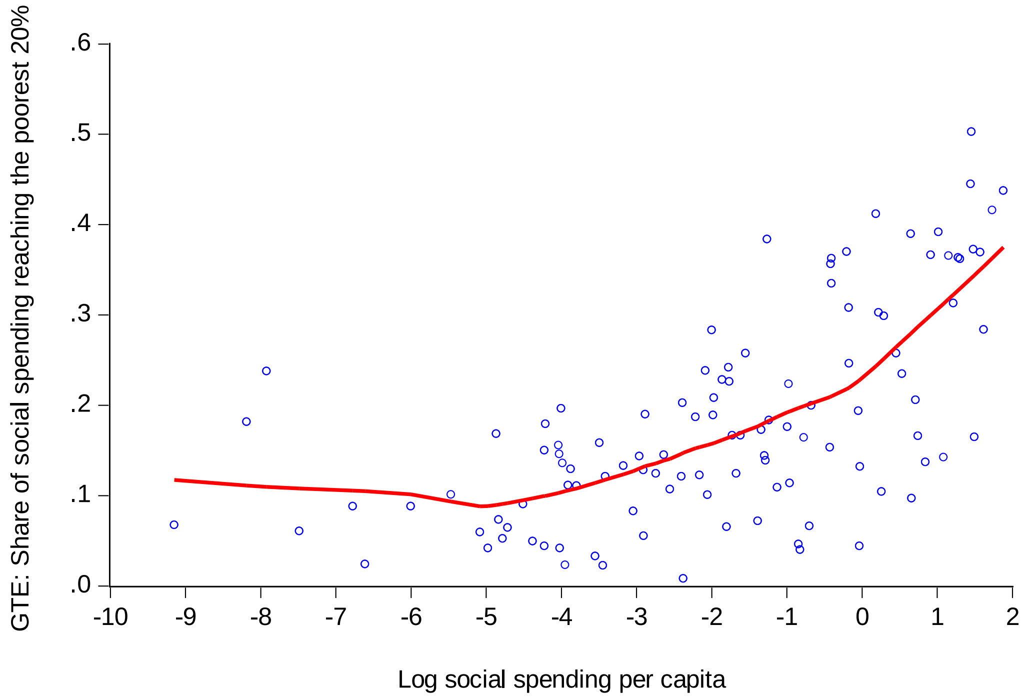

(b) Floor transfer efficiency

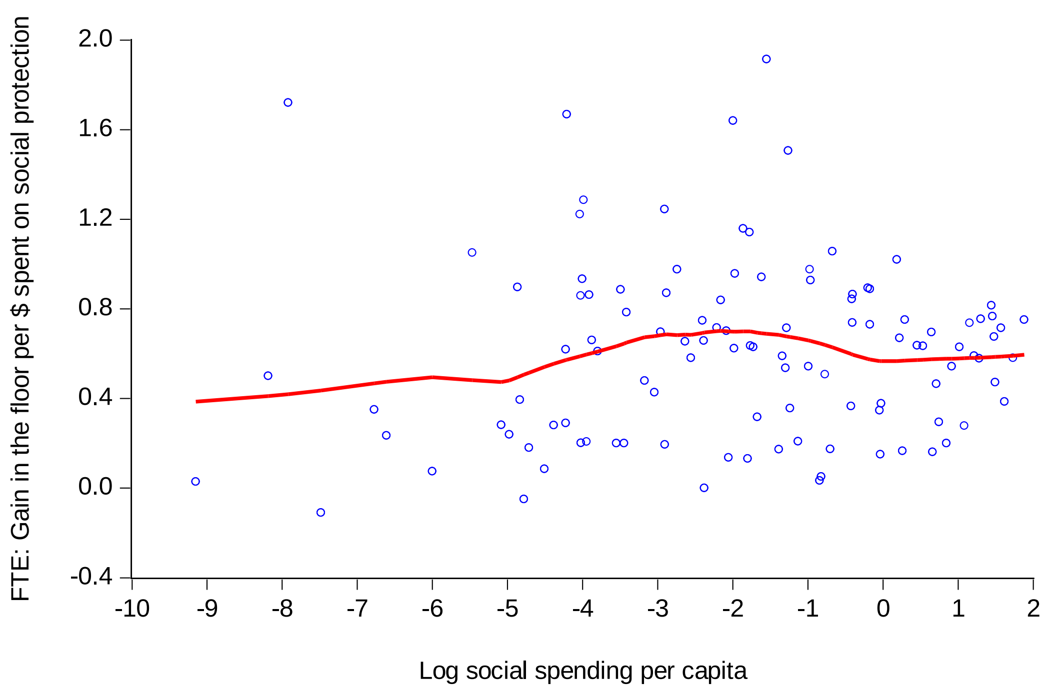

Figure 6: Log median plotted against log mean across countries

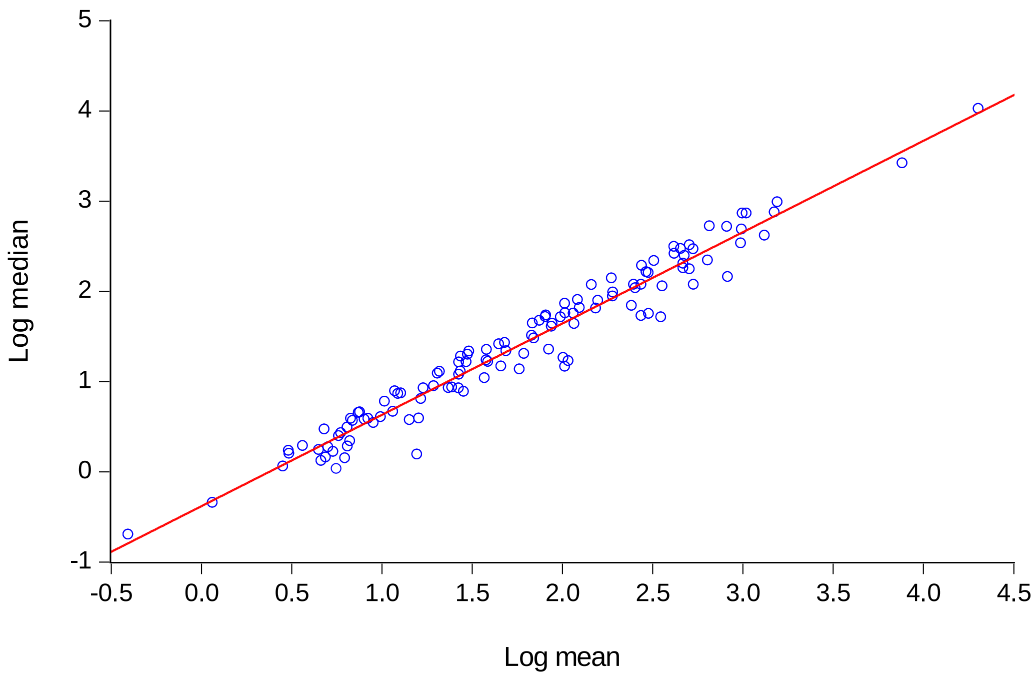

Figure 7: Average transfer plotted against survey median across countries

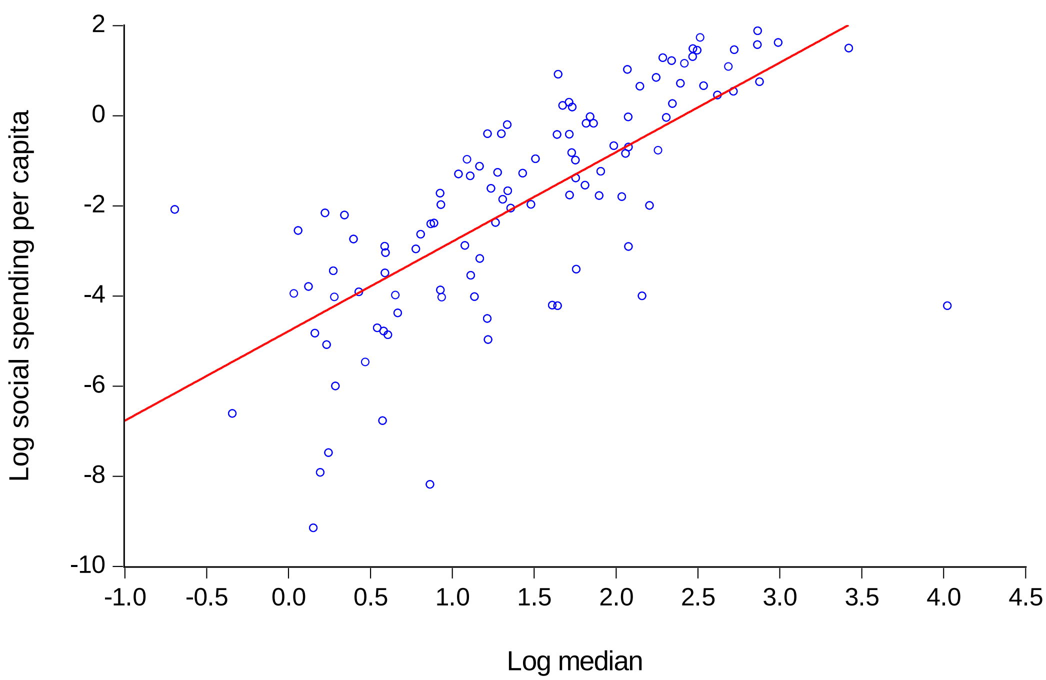

Figure 8: Richer countries have a higher floor

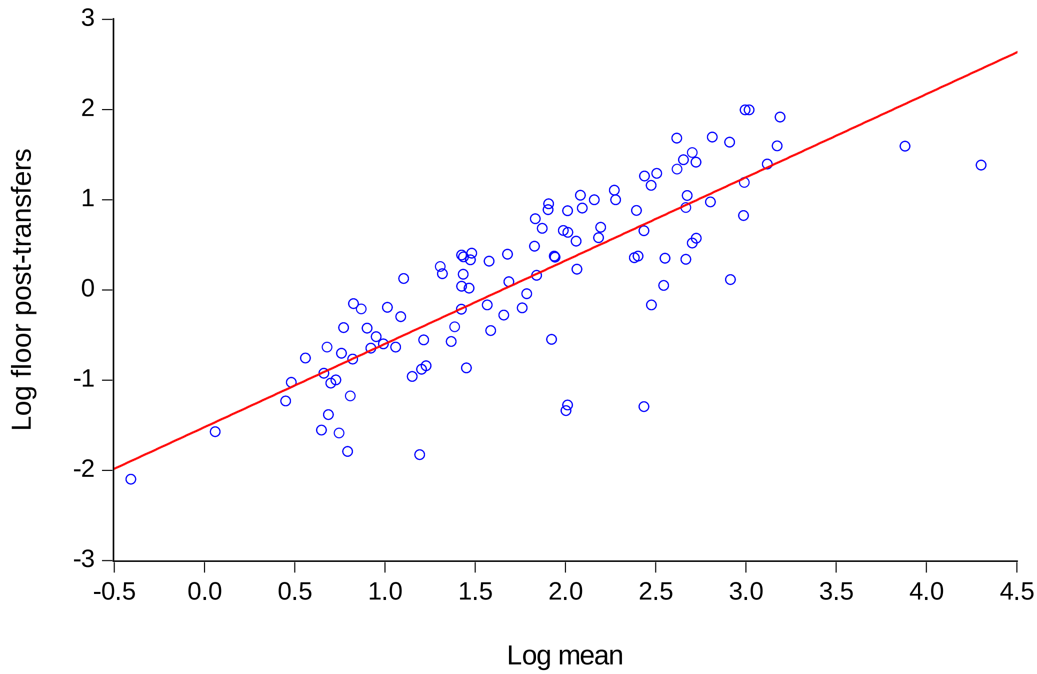

Figure 9: Log pre-transfer floor controlling for the mean plotted against log social spending per capita

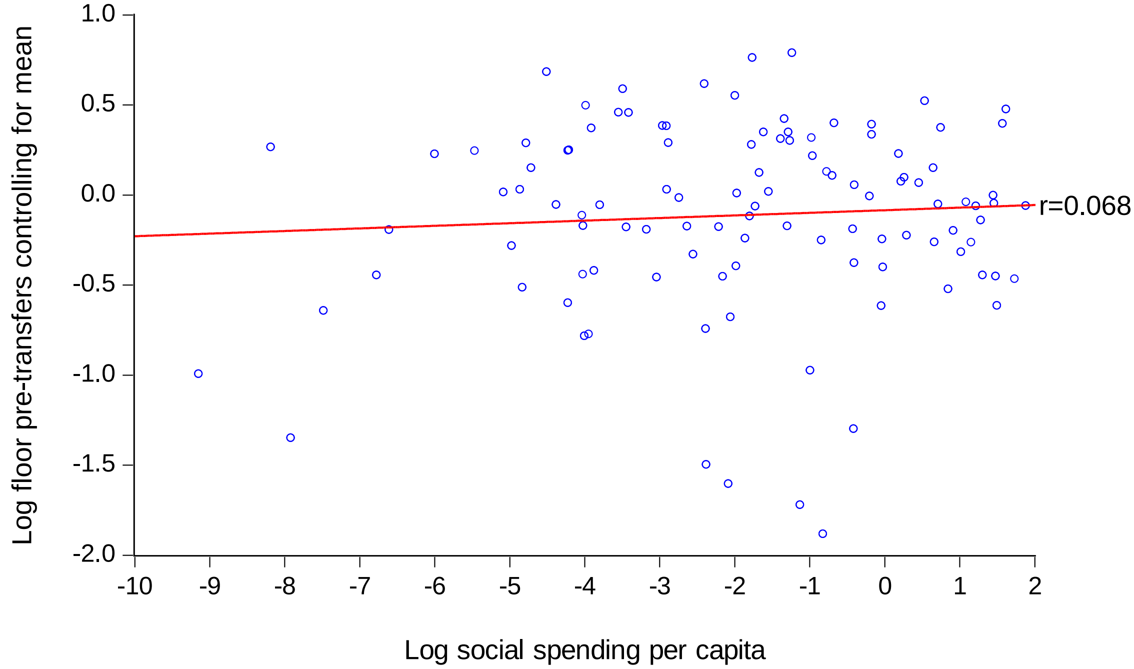

Figure 10: Public spending on SNAP per capita of US population

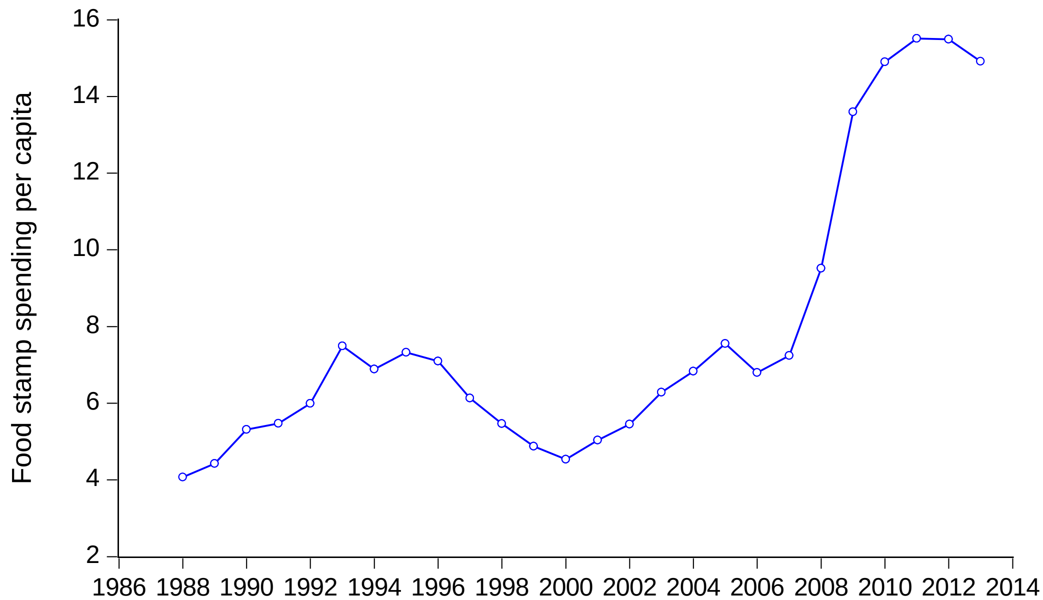

Figure 11: Poverty measures for the US

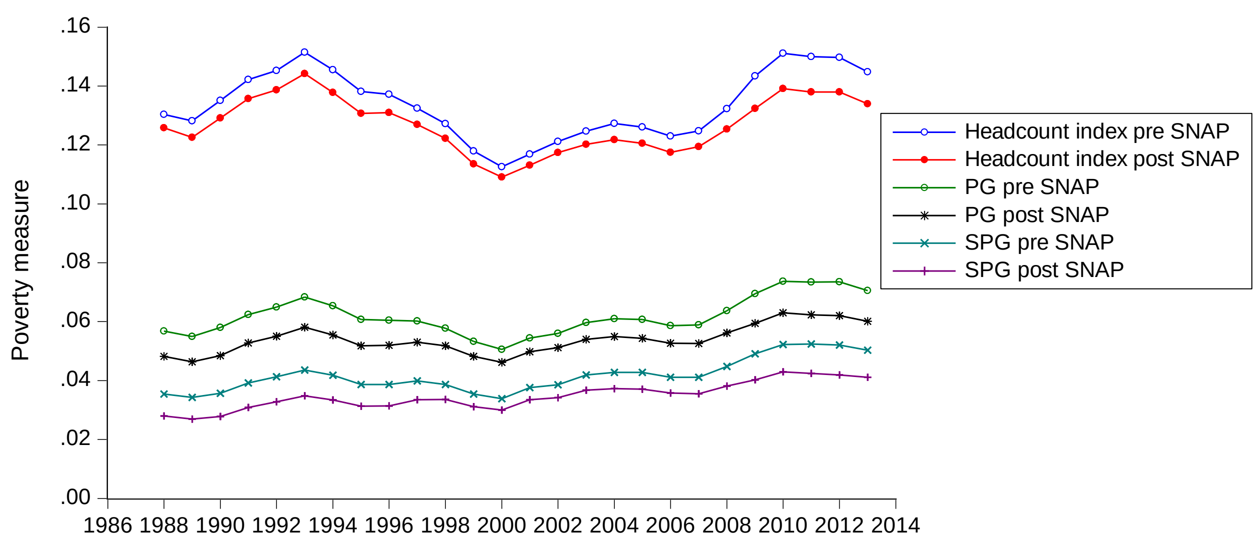

Figure 12: Floor in the United States before and after food stamps

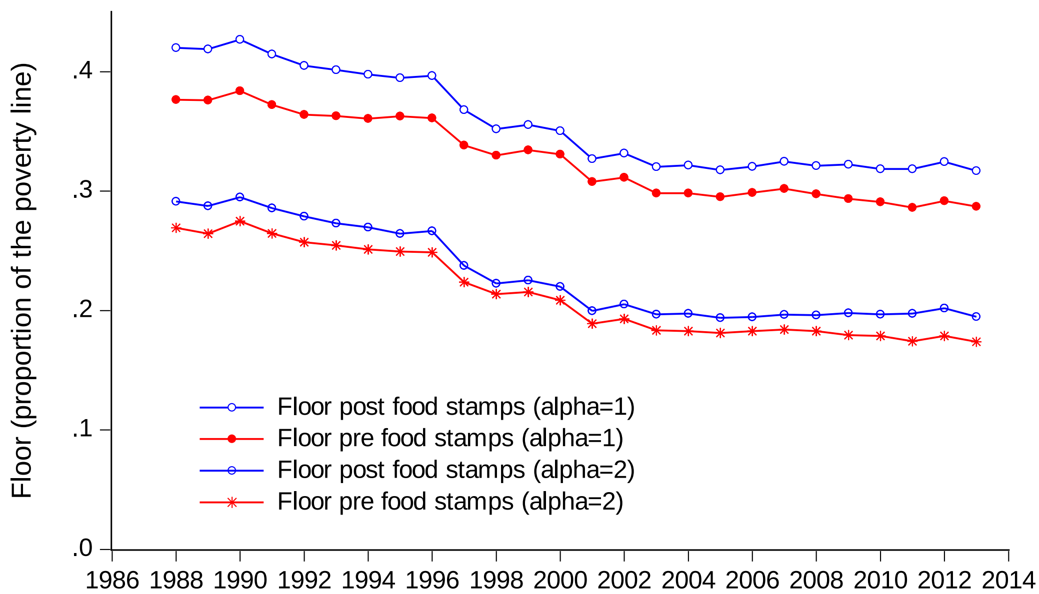

Figure 13: Floor transfer efficiency for SNAP

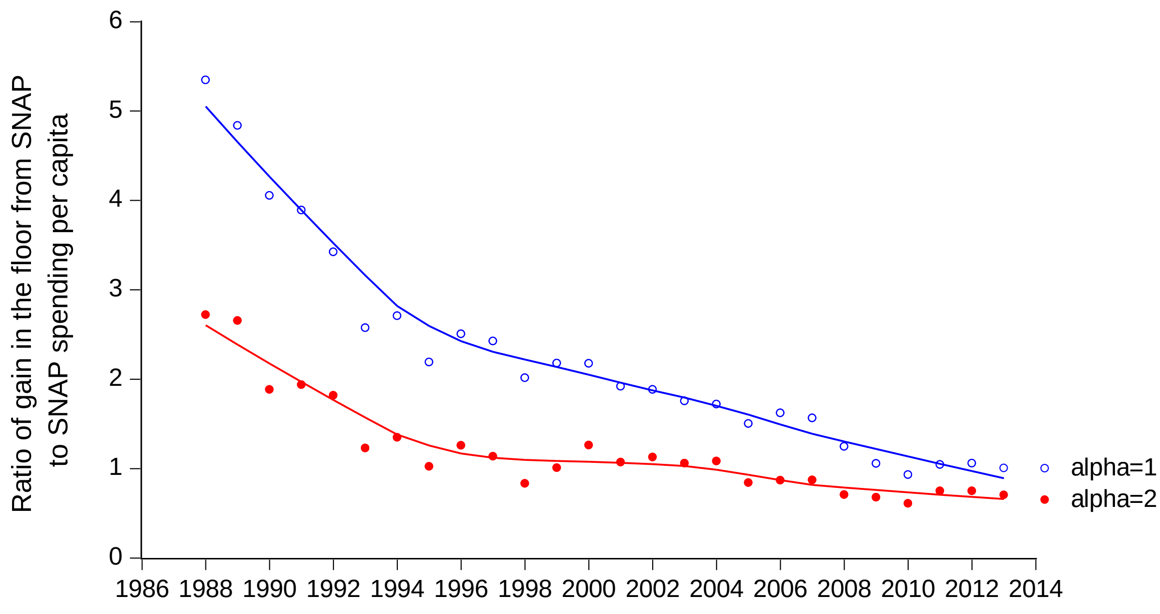

Figure 14: Floor transfer efficiency plotted against SNAP spending per capita

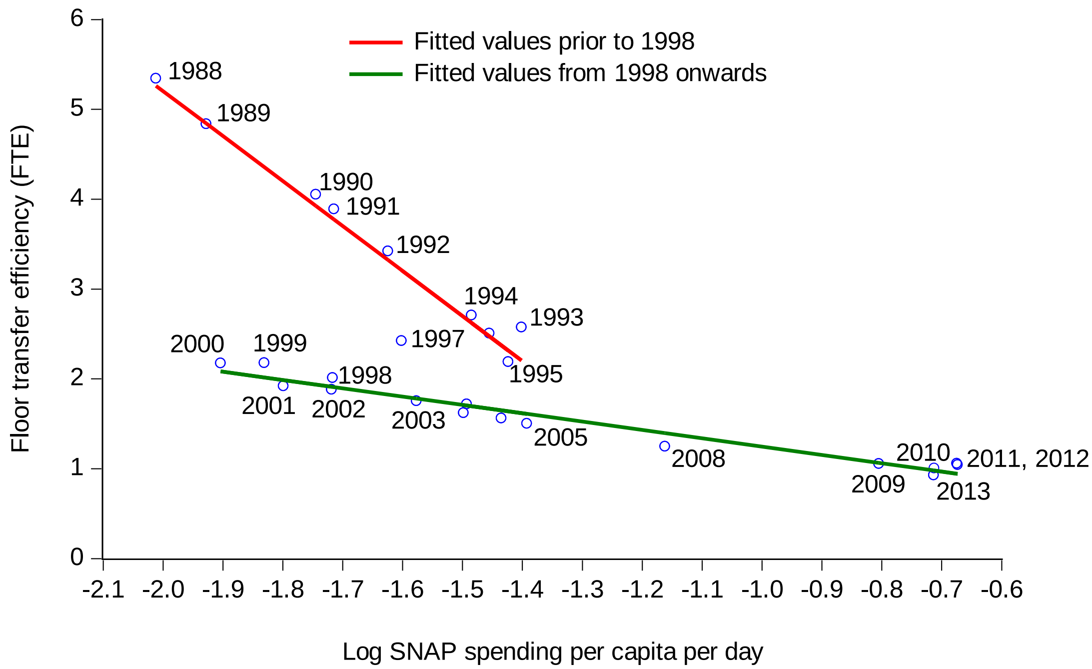

Note: Some year labels dropped to avoid clutter

Figure 15: Simulated floor without the expansion in food stamps in 2000s

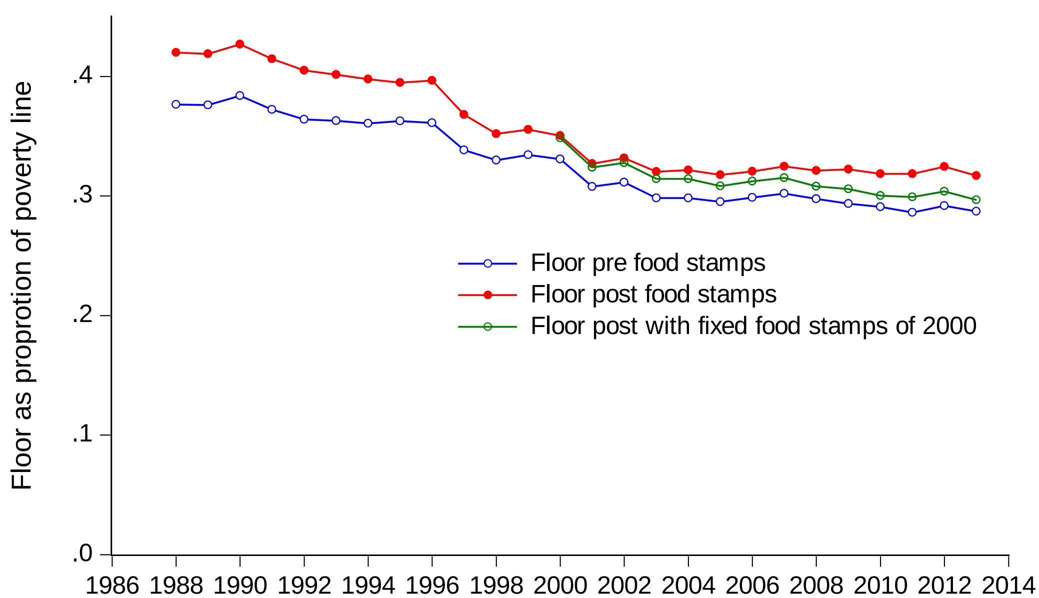

Table 1: Summary statistics for cross-country data set

	Mean	St. dev.	Min	Max
Threshold (z)	3.09	2.99	0.20	17.48
Mean public transfers (τ)	0.88	1.44	0.00	6.56
Floor post transfers ($\hat{y}_{min}^{*post}$)	1.73	1.60	0.12	7.34
Floor post transfers as share of threshold	0.574	0.096	0.228	0.729
Floor pre transfers ($\hat{y}_{min}^{*pre}$)	1.21	0.95	0.03	4.82
Floor pre transfers as share of threshold	0.461	0.139	0.174	0.728
GDP per capita	11.60	11.38	0.57	45.24
Mean transfer to poorest 20%	0.51	0.98	0	5.22
Survey mean (m)	8.41	9.03	0.67	74.05
Survey median (y^{med})	6.01	6.68	0.5	55.97
Headcount index post transfers (%)	26.86	9.55	20.00	56.83
Poverty gap index post transfers (%)	5.83	1.57	3.51	11.19
Poverty gap index pre transfers (%)	10.87	7.26	3.57	36.90

Note: All values displayed above are in daily per capita US$ units, in 2005 prices (at PPP) unless noted otherwise. Public transfers comprise all social insurance, social assistance and labor market programs (classified as “social protection” by the World Bank). N=121, but varies by row depending on data availability.

Table 2: Regression for social spending as a function of the median and mean

	(1)	(2)
lny^{med}	3.831***	0.838**
	(0.958)	(0.413)
lnm	-1.964*	
	(1.022)	
lnGDP		1.247***
		(0.273)
Constant	-3.893	-12.950
	(0.667)	(1.660)
R^2	0.540	0.660

Note: N=110. The dependent variable is $\ln\tau$. Robust standard errors in parentheses. ***: 1% significance;**: 5%; *10%.

Table 3: Regressions for log floor, cross-country data set

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Log floor, pre-transfers		Log floor, post-transfers			Gain in log floor due to social spending			
Log SP transfers per capita ($ln\tau$)	0.027		0.145***		0.106**	0.118***	0.095***	0.058**	0.244***
	(0.033)		(0.025)		(0.044)	(0.018)	(0.012)	(0.027)	(0.043)
Log mean income (lnm)	0.740***	0.792***	0.642***	0.923***	0.686***	-0.099*		-0.030	-0.347**
	(0.105)	(0.065)	(0.070)	(0.062)	(0.084)	(0.052)		(0.066)	(0.155)
Interaction effect ($ln\tau . lnm$)					0.023			0.035**	0.025*
					(0.019)			(0.014)	(0.013)
Log Gap Transfer Efficiency (GTE)									0.664***
									(0.138)
Interaction effect ($ln\tau . lnGTE$)									0.104***
									(0.014)
Interaction effect ($lnm . lnGTE$)									-0.151**
									(0.064)
Constant	-1.399***	-1.544***	-0.738***	-1.530***	-0.843***	0.661***	0.439***	0.497***	1.701***
	(0.249)	(0.131)	(0.160)	(0.109)	(0.200)	(0.128)	(0.039)	(0.152)	(0.340)
R^2	0.629	0.626	0.738	0.665	0.740	0.495	0.462	0.536	0.800
Evaluated at mean points									
$ln\tau$					0.147***			0.121***	0.088***
					(0.025)			(0.020)	(0.016)
lnm					0.642***			-0.098	-0.101**
					(0.070)			(0.064)	(0.051)
$lnGTE$									0.191***
									(0.029)

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

Table 4: Data and summary statistics for the US

	GDP per capita	Mean	Median	Gini index	SNAP per capita	Poverty rate (pre-SNAP)	Poverty rate (post-SNAP)	Floor (pre-SNAP)	Floor (post-SNAP)
1988	94.74	35.86	27.95	0.404	4.07	13.04	12.58	0.377	0.420
1989	97.31	38.23	29.60	0.406	4.42	12.82	12.25	0.376	0.419
1990	98.07	39.22	30.41	0.404	5.31	13.51	12.91	0.384	0.427
1991	96.70	40.04	31.36	0.407	5.48	14.22	13.57	0.372	0.415
1992	98.81	40.95	32.19	0.410	5.99	14.52	13.86	0.364	0.405
1993	100.22	41.89	32.51	0.417	7.49	15.14	14.42	0.363	0.401
1994	103.01	43.90	34.03	0.416	6.89	14.55	13.78	0.361	0.398
1995	104.56	48.21	35.56	0.444	7.33	13.81	13.07	0.363	0.395
1996	107.27	50.65	36.71	0.448	7.10	13.72	13.10	0.361	0.396
1997	110.76	53.72	38.94	0.450	6.13	13.25	12.69	0.338	0.368
1998	114.35	56.37	41.10	0.448	5.46	12.72	12.22	0.330	0.352
1999	118.34	58.34	42.99	0.438	4.87	11.79	11.35	0.334	0.356
2000	121.85	62.35	45.07	0.451	4.53	11.25	10.91	0.331	0.351
2001	121.82	63.83	45.69	0.457	5.03	11.69	11.31	0.308	0.327
2002	122.82	63.82	46.30	0.456	5.45	12.12	11.74	0.311	0.332
2003	125.10	65.20	47.26	0.458	6.29	12.46	12.02	0.298	0.320
2004	128.67	66.49	48.01	0.460	6.83	12.73	12.18	0.298	0.322
2005	131.75	69.95	50.23	0.463	7.56	12.60	12.05	0.295	0.318
2006	133.99	73.73	52.85	0.466	6.80	12.30	11.75	0.299	0.321
2007	135.07	75.24	54.79	0.457	7.24	12.48	11.94	0.302	0.325
2008	133.42	75.71	54.79	0.461	9.52	13.23	12.54	0.298	0.321
2009	128.58	74.76	53.70	0.468	13.60	14.34	13.24	0.294	0.322
2010	130.74	74.89	53.97	0.468	14.90	15.11	13.91	0.291	0.319
2011	131.85	77.74	54.80	0.476	15.51	14.99	13.80	0.286	0.318
2012	133.81	79.61	56.51	0.477	15.49	14.97	13.80	0.292	0.325
2013	135.10	81.39	58.09	0.475	14.92	14.48	13.39	0.287	0.317
Mean	117.64	59.70	43.67	0.446	7.85	13.38	12.71	0.327	0.357
St.dev.	7.65	14.94	9.77	0.025	3.71	1.15	0.94	0.034	0.041

Note: Monetary values in $US per day except SNAP is per month. Gini index is for families (as usually calculated for the US). Poverty rate as % of population. Floor as a proportion of the threshold. Authors' calculations of means, poverty and inequality measures based on CPS micro data.

Table 5: Regressions for the US floor before and after SNAP

	(1) Log floor pre-SNAP	(2) Log floor post- SNAP	(3) Gain in log floor due to SNAP
Lagged dep.var.	0.527***	0.490***	0.433***
	(0.137)	(0.137)	(0.154)
$ln\tau$	0.003	0.028**	0.028***
	(0.013)	(0.012)	(0.009)
lnm	-0.184***	-0.233***	-0.039***
	(0.051)	(0.066)	(0.014)
Constant	0.838***	1.151***	0.283***
	(0.244)	(0.346)	(0.096)
R^2	0.963	0.967	0.879
DW	2.091	2.091	1.676

Note: N=26. The dependent variable is $\ln\hat{y}_{min}^{*post}$. HAC standard errors in parentheses. ***: 1% significance;**: 5%; *10%.

Table 6: Augmented regressions for the proportionate gain in the floor attributed to SNAP

	(1)	(2)
$ln\tau$	-0.474***	-1.789***
	(0.059)	(0.184)
lnm	-0.195***	
	(0.015)	
$lnGDP$		-0.450***
		(0.037)
$ln\tau.lnm$	0.068***	
	(0.008)	
$ln\tau.lnGDP$		0.170***
		(0.017)
Constant	1.477***	4.825***
	(0.113)	(0.386)
R^2	0.919	0.897
DW	1.491	1.272

Note: N=26. The dependent variable is $\ln\left(\frac{\hat{y}_{min}^{*post}}{\hat{y}_{min}^{*pre}}\right)$. HAC standard errors in parentheses. ***: 1% significance;**: 5%; *10%.