WIDER Development Conference

Public economics for development

5-6 July 2017 | Maputo, Mozambique

A dose-response function approach for labour supply and cash transfers. The case of Zambia

Abstract

Cash transfer programs pursue mainly protective objectives, but can also impact rural livelihoods by inducing investments in productive activities and changing household labor allocation. We adopt a continuous treatment approach to quantify how households' labor supply responds to transfer size. We find a shift from paid labor to own farm labor and find that the transfer size is well within a level that would have disincentive effects on time spent on own farm activities. The switch from paid to own farm labor occurs at lower levels of transfers for labor-constrained households, and at higher levels for non-labor constrained households.

Keywords: cash transfers; labor supply; continuous treatment approach, IPW-regression

JEL Classification: C21, J22, H24

Introduction

Despite Zambia's strong economic recovery and robust growth over the past decade impacts on poverty rates have been limited. In response to this, Zambia has followed other countries in Africa in launching unconditional cash transfer (CT) programs as part of social protection systems, with the aim of reducing poverty and hunger. Although poverty alleviation is the main goal, these programs may also have significant impacts on rural livelihoods by inducing investments in productive activities and increased production. The literature on CTs has focused mainly on the evaluation of their primary intended outcomes, such as consumption, food security, education and poverty, and less on second-order effects on income generating activities that might be triggered by the injection of liquidity in the household's economy (Rawlings and Rubio, 2003). Since most effects in productive activities pass through labor decisions, in this paper we focus on how the amount of the transfers affects adults' labor supply and work incentives of recipient households.

Supporters of CTs see the transfers as a tool to give those at the very bottom of the income ladder a leg up. On the downside, there is increasing concern among policy-makers of the possible incidence of CT on work incentives. In fact, critics of these programs argue that handouts may stifle economic activity and create disincentives to work and a culture of dependency. Recent studies focusing on labor outcomes of CTs in sub-Saharan Africa do not show consistent evidence in this regard. In their analysis of Kenya's CT-OVC, Asfaw et al. (2014) did not find significant impacts on labor reallocation overall, even though some heterogeneous impacts by age and gender were observed. Gilligan et al. (2009) show no disincentive effects on labor supply looking at Ethiopia's Productive Safety Net Program (PSNP), while Ardington et al. (2009) find positive impacts on adult labor supply for South African Old Age Pension.

In Latin America, conditional CTs do not appear to have much impact on work incentives and adult labor supply. Studies from Brazil (Foguel and Barrios, 2010; Ribas and Veras Soares, 2011), PROGRESA in Mexico (Skoufias and Di Maro, 2008) and Honduras (Alzúa et al., 2013), using a variety of approaches, did not find a significant impact on participation in adult wage employment. Some evidence shows, however, that CCTs may modestly reduce the time spent working, for males in Nicaragua and females in Brazil, as well as substitution between wage employment and domestic housework in Brazil.

Economic theory suggests several ways of how unconditional CTs might affect adult labor supply in recipient households. CTs constitute an increase in non-labor income which causes a loosening of the household budget constraint making work less attractive relative to leisure. However, in the presence binding credit and liquidity constraints selling their own labor might be the only viable strategy for adult household members to obtain some liquidity and meet their consumption or investment needs (Rose, 2001). This can lead to distortions in livelihood strategies as when households are forced to sell more labor off-farm than would be optimal. In this respect, unconditional CTs can provide a secure source of liquidity that may help rural households overcome credit market failures and consequently allow them to work more on their own farms thus affecting off-farm and on-farm labor allocative efficiency. The interplay of these channels makes it an empirical question whether, and to what extent, a given amount of CT affects the labor supply and work incentives of recipient households.

In this paper we are interested in the impacts of transfer size on labor supply. This is obtained by comparing the average labor supply of those exposed to a given level of treatment with the average outcome of the controls. This question is relevant from both a policy perspective and an empirical point of view. On the one hand, governments are interested in knowing not only whether providing cash to poor vulnerable households create disincentives to work or not, but also whether it exists a threshold level in the size of the transfer below which transfers can contribute to developmental objectives without creating welfare traps and distorting the domestic labor market. Further, governments in many low-income countries are worried about the long-term fiscal sustainability that is implied when the cash transfer is offered above a certain amount. This concern, however, is not tackled in the paper. On the other hand, by adopting a continuous treatment approach, we are able to identify causal pathways between changes in labor supply and the magnitude of the change in unearned income given by the amount of the received transfer. The main econometric problem is that treatment was not randomized among the different transfer amounts which opens the way to potential endogeneity bias. We adopt two estimation strategies for the continuous treatment effects: the first one is a doubly robust approach that combines OLS regression with inverse probability of treatment reweighting and relies on a selection-on-observables kind of identifying assumption; the second one recognizes that the treatment might be endogenous allowing for selection on unobservables and consists of an instrumental variables approach.

To the best of our knowledge this is the first empirical paper that studies how household labor response changes with the amount of the received cash transfer. This study is the natural continuation of a previous paper from the same authors that focused on the binary treatment effect of the cash transfer program obtained by the comparison of the treated as a whole, regardless of the transfer size, and the controls (Prifti et al., 2016).

Cash Transfers and labor supply: conceptual framework

Here we illustrate some standard predictions on the relationship between the amount of cash transfers, considered here as a shock to unearned non-labor income, and the labor supply of rural households. The formal model based on the traditional income-leisure approach is presented in Del Carpio (2008) while here we focus on the ideas behind it illustrated in figure 1. The upper part of the picture shows the non-labor income elasticity of labor supply that can be estimated with a reduced form regression of labor on the transfer amount, while the lower part refers to the relationship between levels of unearned income and supplied labor which can be derived from the former. The relationship between changes in non-labor income as given by the CT and changes in supplied labor is monotonically declining, thus the curve crosses the horizontal axis at some critical level of the CT. In fact, a non-constant income elasticity implies a nonlinear relationship between labor supply and non-labor income. More specifically, a decreasing trend in the elasticity implies a concave shape of the relationship between labor supply and non-labor income. Finally, the alternation of sign in the elasticity curve hints at a non-monotonic relationship of labor supply and unearned income. This implies an inverted-U shape relationship between the levels of non-labor income and levels of labor supply that peaks at the abovementioned critical level of the transfer size. The inverse U shape reflects the trade-off between a higher utility from more consumption and the diminished utility due to the leisure that has to be given up to increase labor and production. In the increasing portion, the income effect of increased consumption prevails on substitution effects induced by the desire to substitute work for leisure. The opposite holds true in the declining part of the parabola. Which of these two opposing forces will prevail in the trade-off becomes an empirical question. Predictions on both relationships find significant support at the empirical level in our study.

The theoretical setup in previous studies on this topic has imposed a constant effect of non-labor income on labor supply estimated with a dummy treatment approach. Conversely, in this paper, we allow the effect to vary with the level of unearned income and adopt an empirical strategy that allows us to estimate curve wise instead of point estimates of the effects.

In this conceptual framework we also differentiate between the two types of labor we take in consideration in our study, namely paid labor and on-farm labor[1]. Assuming that disutility from selling labor to other villagers is higher than the disutility of own farm labor and that the marginal productivity of own farm labor is higher than marginal productivity of paid labor it can be shown that as non-labor income increases, paid labor increases more slowly or might even decrease compared to own farm labor, leading to a reallocation of household labor from the former type to the latter.

Program design, data and descriptive statistics

The Government of Zambia started implementing the Social Cash Transfer (SCT) in 2003 and gradually expanded its coverage over time. By 2015, the SCT program was implementing four

[1] Paid labor and wage labor are used interchangeably in this study and both refer to off-farm labor for which the adult receives a payment in cash.

different targeting models: labor-constrained, universal old-age pension, child grant and multiple categorical models. The Child Grant (CG) model of the SCT started in 2010, aiming at alleviating poverty among the poorest households and blocking its intergenerational transmission. The pilot evaluation of the CG model was implemented in three districts that had never received any CTs and with highest rates of mortality, morbidity and stunting among children under 5 years of age. The CG was based on a categorical targeting mechanism, reaching any household with a child under 5. Beneficiary households received 60 new kwacha (ZMK) a month. The planned transfer size was constant regardless of the number of household members and corresponded to about 28 percent of the median household's monthly consumption expenditure. Payments were unconditional. The designated recipient of the cash was the female head of household. Less than 10 percent of CG recipients ever reported having to make multiple trips to receive a single payment and only 2 percent of them reported of missing any payment (American Institutes for Research, 2013).

CG impact evaluation was designed as a longitudinal randomized controlled trial (RCT) with random assignment at the community level. Random assignment of the communities to treated and control groups occurred only after baseline data were collected, thus avoiding anticipation effects in the baseline data. The final baseline sample is composed of 2,515 households. Baseline data were collected during the lean season that spans from September through February, during which people have little food left from the previous harvest and hunger is most felt. The 24-month follow-up data collection occurred in September and October 2012 exactly 24 months from the baseline study thus avoiding seasonal effects.

In this paper we are interested in comparing the average labor supply of those households exposed to a given level of treatment with the average outcome of the controls. The binary randomization mechanism of the RCT design should ensure comparability along every observed and unobserved dimension between treated as a whole, regardless of treatment level, and the controls. However, randomization does not operate with regard to the different doses of treatment. In this sense, the treatment level modeled by the self-declared amount of cash a household received during the 12 months preceding follow-up data collection deserves special attention. Establishing where the identifying variation in the treatment variable is coming from is important in order to attribute the observed effects to the different amount of cash transfers that exogenously increase household non-labor income and not to other observed or unobserved determinants of household labor supply. In Zambia the amount of cash households receive was supposed to be flat regardless of household size and of the number of children under 5. Although the theoretical yearly amount of the transfer before the follow-up survey was supposed to be 720 new ZMK, corresponding to six payments of 120 new ZMK, in the course of the evaluation few households received an additional seventh payment. This means that the theoretical upper limit of the transfer may reach 840 new ZMK for some households. However, there is some variability in the amount of cash received by the households during the last year. Figure 2 shows the histogram of the amount of cash received by the treated households during the last year the program was in place. The average yearly amount of self-reported transfer by beneficiaries equals 639.5 new ZMK with a standard deviation of 170 new ZMK. As with any self-declared variable there could be issues of misreporting. Some households reported having received more than the planned amount and others less. From an econometric point of view misreporting translates into measurement error in the continuous treatment variable which leads to endogeneity of the latter and to biased coefficient estimates. Figure 2 lends some support to the misreporting interpretation of the causes of variation of the continuous treatment. In fact the superimposed Gaussian-looking distribution may indicate noise in the transfer values due to misreporting. Another cause for endogenous variation in the supposedly flat transfer size may be related to missed payments, although this is less likely since only a small fraction of the beneficiaries report to have missed any payments. Endogeneity in this case would arise because the factors leading a household to miss a payment, such as lack of time or able-bodied members, lack of information and low social capital may all influence labor supply.

Therefore, both possible sources of variation of the transfer cannot be reasonably assumed to be exogenous and may therefore lead to endogeneity issues. Whatever the source of endogeneity in the continuous treatment it raises the issue of correcting the resulting potential bias. We address this issue in the following sections.

Table 1 shows descriptive statistics by treatment arm. The treatment and the control groups are observationally equivalent in terms of almost all included confounders. We do control for the observed differences by including those confounders as controls in the regression analysis in the results section. Finally, figure 3 gives an illustration of the relationship between our main outcome variables of interest – total days per week of paid and on-farm work by all adult members – and the continuous treatment. We group the sample based on the quintiles of the positive part of the treatment variable and plot the average outcome for each quintile (columns 1-5) and for the controls (column 0). The left-hand side plot shows that average paid labor for the treated is lower than for the controls across all quintiles, but does not follow any discernible pattern. For own farm labor we observe that the average outcome for the treated is higher compared to the controls across all quintiles and follows a clear inverse U-shaped pattern. The supplied quantity of on-farm labor by the households increases with the size of the transfer up to a certain level of treatment close to the average size and then declines. These two descriptive graphs based on the raw average of the outcomes at different treatment levels give only an approximate idea of the relationship between transfer size and labor supply as the effects of the former may be confounded by other observed and unobserved factors that were not taken into account here. Below, we control for these factors in order to shed light on the causal relationship between transfer size and labor supply.

Empirical approach

The setup for the continuous treatment approach adopted in this paper builds on Cerulli (2014), Robins et al. (2000) and Uysal (2015). The treatment t_i takes values in some general interval J and $D_j(t_i)$ is an indicator for individual i receiving treatment level j, i.e. $D_j(t_i) = 1$ if $t_i = j$ and 0 otherwise. Let also w_i=**1**[t_i>0] be the derived binary treatment indicator that equals one for the group of treated subjects as a whole. Our causal parameter of interest is the average treatment effect on the treated for a particular level of treatment t i.e., $ATT(t) = E[Y_i(j) - Y_i(0)|t_i = j]$. This is also known as the *relative dose response function* (DRF) which is shown to identify the average difference in outcomes between a group of treated subjects that received a given dose of treatment and what would have been observed had they not received any treatment. Estimation of the causal parameters of interest is based on the following linear regression model (Cerulli, 2014)

$$Y_i = \delta_0 X_i + (\delta\bar{X} + \bar{h})w + \delta(X_i - \bar{X})w_i + \big(h(t_i) - \bar{h}\,\big)w_i + \varepsilon_i \qquad (1)$$

where t_i is the level of CT, $h(t_i)$ is some general function of the continuous treatment level and $E[\varepsilon_i] = 0$. In the continuous treatment setup the average difference in the outcome of the treated and the controls, given their observed characteristics, is allowed to vary by an arbitrary function of the treatment level $h(t_i)$. The average causal effect at a certain level of treatment is estimated as $ATT(t) = ATT + (h(t)_{t>0} - \bar{h}(t)_{t>0})$. We assume a cubic polynomial for $h(t)$= at^3+bt^2+ct +d whose parameters are estimated along with the rest of the coefficients in equation (1). The parameters are estimated as $\boldsymbol{\beta}_{OLS}$=$(\boldsymbol{XX})^{-1}\boldsymbol{X}'y$ where $\boldsymbol{X}$={$X, w, h(t), \bar{X}, \bar{h}$}, $\boldsymbol{\beta}_{OLS}$={$\delta_0, \delta, a, b, c, d$}.

We focus on the intensive margin by taking Y to be the total number of days per week worked by all adult household members. Equation (1) is estimated separately for the two main types of labor supply we consider in this study, namely, any wage labor and own farm labor and for a third variable, hired labor for agricultural use. The vector of observed pre-treatment characteristics X_i includes household level controls listed in table 1.

As to the issue of the identifying assumption, the main issue is that households were not randomly assigned across the different levels $t \in \mathcal{J}$, therefore we need to make extra assumptions to assure consistent estimation of the *ATT*(*t*). In particular, estimation of the continuous treatment parameter proceeds in two steps: first we assume conditional mean independence (CMIA), i.e. that conditional on observed pre-treatment characteristics households are as good as randomly assigned to a given transfer amount (*selection on observables*). We formally discuss this assumption in the next subsection and provide evidence for its validity in the results section; second we relax the conditional independence assumption and allow for selection on unobservables in an instrumental variable setup (*selection on unobservables*).

Selection on observables

Unbiased estimation of the treatment effect for a certain level of treatment (*ATT*(*t*)) in the outcome model expressed by equation (1) requires that the functional form assumption and the identifying assumption, namely linearity and the CMIA, are valid for our data. Formally, the CMIA states that $E[Y_i(j)|X_i]$= $E[Y_i(j)|D_j(t_i) = 1, X_i] =$ $E[Y_i|D_j(t_i) = 1, X_i]$= $E[Y_i|t_i = j, X_i]$. As a result, any observed differences in the outcome of interest between those exposed to some level of treatment and the controls are attributable only to the treatment.

To gain some protection against bias from the possible failure of these assumptions we adopt a doubly robust estimator in the form of regression with Inverse Probability Weighting (IPW). First, doubly robust (DR) estimators may mitigate selection bias arising from non-random treatment assignment which motivates in part its use in the current context. Secondly, the assumed linear regression model for the outcome is probably a rough approximation of the true model relating the outcome variable Y to t and X. In addition to the outcome model, the doubly robust approach envisages a second model for the probability of treatment relating t to X, i.e. the generalized propensity score (GPS). The distinguishing feature of DR models is that coefficient estimates in the outcome model are consistent if either model is correctly specified in terms of the functional form and relevance of the included covariates (Kang and Schafer, 2007). We use stabilized weights for the DR approach given by $S\varpi = \frac{P(t)}{P(t|X)}$where $P(t|X)$ is the GPS and $P(t)$ is the marginal density of treatment. Our final estimating equation is based on Weighted Least Squares (WLS) version of equation (1). Hence, the parameters are obtained as $\boldsymbol{\beta}_{WLS}$=$(\boldsymbol{X}\Delta\boldsymbol{X})^{-1}\boldsymbol{X}'\Delta y$ where $\boldsymbol{X}$=$\{X, w, h(t), \bar{X}, \bar{h}\}$, $\boldsymbol{\beta}$=$\{\delta_0, \delta, a, b, c, d\}$ and Δ is the weighting matrix featuring $S\varpi$ in the diagonal.

When the GPS is *correctly specified* in terms of functional form and relevance of the included covariates, weighting by $S\varpi$ creates a pseudo population in which X no longer predicts t i.e. $X \perp \mathbf{1}[t = j] \mid P(t|X)$ (balancing property of the GPS). Therefore, accessing whether the balancing property is verified in the data is a fundamental *diagnostic check* for the functional form of the GPS. Consequently, knowing that the GPS model is even nearly correctly specified allows the analyst to claim the double robustness property of removing bias from the possible misspecification of the other model, namely the outcome model (equation 1). Moreover, the balancing property allows also to informally gauge the validity of the CMIA. Since the CMIA implies that subjects exposed to different treatment levels should be equivalent on average both in terms observed and unobserved characteristics except for the level of treatment intake, we can test it by assuring that the observed characteristics are (mean) independent from treatment status. In the context of continuous treatment there are a few statistical procedures that can be applied to assess the validity of the balancing property. We apply two of them. In the first one (Imai and Dyk, 2004), we run regressions of each covariate on the treatment and the GPS. If the covariate is balanced, then the treatment variable should have no predictive power conditional on the GPS. A comparison of this coefficient to the corresponding coefficient of a regression that does not include the GPS can be used to gauge the balancing provided by the GPS. The second approach from Fong and Imai (2014) quantifies imbalance as the sum of the squared Pearson correlations between the treatment and covariates in the original and in the reweighted sample.

Selection on unobservables

Reweighted estimation can only ensure against the selection-on-observables bias but, evidently, provides little relief against the selection-on-unobservables bias. We mentioned that variation in the continuous treatment may be endogenous due to either misreporting or missed payments. An obvious way to address endogeneity bias, regardless of the source, is through instrumental variables, assuming appropriate instruments are available, i.e. variables that can be reasonably assumed to be unrelated to household labor supply and that are strong determinants of the received transfer amount. We follow Fichera et al. (2016) and use two separate instruments in a two-stage residual inclusion (2SRI) model, an alternative implementation of the two-stage least squares model that is consistent in both linear and non-linear models (Fichera et al., 2016).[2] The first stage is simply given by $t = \theta W + e$, where W includes all exogenous regressors in equation (1) and two instruments in two separate models: *i*) the time needed to collect the transfer which includes the travel time from home to the payment point and back again and the time the collector had to wait at the payment point *ii*) the average transfer size at

[2] The 2SRI procedure consists in obtaining the residuals from the first stage and adding them in the second stage regression keeping the original endogenous variable. We prefer this method to the classical two-stage predictor substitution (2SPS) that consists in substituting the endogenous variable with its fitted values from the first stage for comparability reasons. Since the DRF curves obtained by OLS are fitted against the original transfer values we use them in the IV procedure too and add the first stage residuals.

the community level. These are relevant instruments because they are correlated with the treatment amount in the first stage model. However, we note that both instruments have advantages and disadvantages relating to the amount of variation and the potential for indirect pathways to labor supply. On the one hand, the travel and waiting time needed to collect the transfer may clearly explain missed payments but it is hard to see an intuitive and direct correlation with measurement errors. On the other hand, there is less variation in the treatment when using community-averaged values, however a fair amount of variation is left as the standard deviation of the community-averaged transfer size is half that of the original variable. It is unlikely that the community average of the transfer size directly affects the household's engagement in paid or farm labor. Moreover, by using community-averaged values of the continuous treatment variable we can average out part of the measurement error, assuming it has a classical form (Fields et al., 2007). It should be a valid instrument and alleviate concerns of reverse causality and measurement error because the community-averaged transfer should be uncorrelated with household i's unobserved determinants of labor supply and less prone to mismeasurement. The second stage is given by equation (1) with the addition of the estimated residual $\hat{e}$ from the first stage. We have bootstrapped the standard errors with 250 replications. Large divergences among the DRF curves instrumented by either measure could indicate whether any of the above mentioned caveats are a serious concern for our identification strategy.

Results

In this section, we show estimates of the average treatment effect on the treated at treatment level t (ATT(t)) from equation (1). Figure 4 shows the relative Dose Response Function graphs for the number of days per week worked by adult household members in any paid job (left hand side) and in own farm labor (right hand side).

The left hand side of figure 4 shows that the effects on the total worked days in any paid job is negative and lies below the horizontal axis over most of the range of treatment amounts. However, the interval of treatment values over which negative effects are statistically different from zero are mostly concentrated around treatment amount and between 60 and 150 ZMK. The reduction in paid labor is not surprising since wage labor in many African rural contexts refers to casual labor, which often takes the form of piecework weeding or ridging on the fields of fellow smallholders. This is seen as a means of last resort to obtain cash in moments of need, and from an economic theory point of view can be seen as a coping strategy rather than the result of income-maximizing behavior (Fink et al., 2014). Consequently, as soon as families get the much needed cash they pull out of occasional paid labor on other villagers' farms. The CT in this case is providing an external remedy to a failure in the local credit market.

The right hand side of figure 4 shows that the effect increases with larger transfer levels. This increase in labor supply on farm can be interpreted as time spent on casual wage labor being shifted on farm. The statistically significant points of the ATT(t) curve are densely concentrated around the theoretical yearly transfer between 650 and 950 ZMK. We notice that the DRF on the right is the specular image with respect to the horizontal axis of the DRF on the left, both in terms of shape of the curve and statistically significant portions of it. Comparing the significant portion of the curves for off farm wage labor and on farm labor, it appears as though it is the same households that are switching from off farm labor to on farm activities.

The significant increase of on-farm labor for medium to high values of the transfers could be related to the cost of required investments in complementary inputs (purchases of fertilizers and hiring of machinery) necessary to boost farm activities. For treatment levels above 1250 ZMK the effects become negative: households that received more than this amount reduced the supply of on-farm. Thus the potential disincentive effect of cash for on farm activities kicks in at levels well above the CG transfer level.

The estimated relationship between changes in non-labor income, given by the cash transfer, and changes in supplied labor is shown in figure 5 for off-farm paid labor (left hand graph) and own farm labor (right hand graph). The amount of households' off-farm labor supply is monotonically decreasing with the level of unearned income (the cash transfer). In other words, households that receive a greater transfer tend to sell less labor compared to households who receive a smaller transfer. While this might seem reasonable, it provides only partial support to our economic model from the second section, as we are capturing only the decreasing arm of the curve.

On the other hand, the relationship between own farm labor and transfer size has an inverted-U shape, thus following the theoretical predictions. On-farm labor supply increases with non-labor income and peaks at the threshold level ($t^{*}\approx$ 1250 ZMK) mentioned above corresponding to the treatment level where the relative DRF curve crosses the horizontal axes (ATT(t^{*})=0). In terms of our conceptual framework, households increase the amount of supplied labor until the utility from the resulting increased production and consumption outweigh the disutility from forgone leisure. For levels of unearned income above the threshold level the opposite stands true and households begin to decrease on-farm labor supply accordingly. The reason of the partial misalignment of empirical findings on the off-farm labor supply from the theoretical predictions may lie in the fact that while the amount of on-farm labor is the result of income- and utility-maximizing behavior, off-farm labor results from the absence of alternatives to meeting temporary consumption and liquidity needs. Unfortunately we cannot determine what exactly households devote time to (whether leisure, household chores, or child care) because we do not have information on time use. We leave this issue open to future research.

Finally, we look at the effects of the cash transfer on the days per week of labor hired in for agricultural production. The picture is omitted to save on space. The pattern for hired labor is very similar to on-farm own labor supply. Treated households hire more labor for agricultural use compared to the control group regardless of the treatment level. The increase is considerably higher than the increase in the household's on-farm own labor supply. This result can be interpreted in the same fashion as the on-farm labor, being related to the cost of complementary inputs.

Robustness checks

We now turn to diagnostics of our benchmark model consisting of a doubly robust regression and provide evidence in support of the adequacy of the choice of the functional form and of its main identifying assumption, respectively, in the GPS model and in the outcome model. In order to gauge the covariate balancing property of the estimated GPS we first run two sets of regressions of each covariate. In the first set, only the treatment is included in the right-hand side, while in the second set we include the treatment and the estimated GPS. We follow a standard approach in the literature and summarize the results of this balancing check using normal quantile (Flores et al., 2012). The left-hand side graph of figure 7 shows the standard normal quantile plot for the t-statistics on the regressions that do not include the GPS. In the full sample there are 4 covariates with t-statistics greater than or near 2 in absolute value. This gives an idea about how unequally spread the covariates are across groups when not controlling for the GPS. The right-hand side graph in figure 7 shows that, once the GPS is included in the regressions, there are no more statistically significant t-statistics, that is, the balance of the covariates improves. As a second balancing check we follow the approach from Fong and Imai (2014). We compute pairwise Pearson correlations of each covariate in the outcome model with the treatment variable. Although correlations are low and mostly not statistically different from zero there are three of them that are statistically significant. Correlations in the reweighted sample are all statistically insignificant. Moreover, quantifying imbalance as the sum of the squared Pearson correlations between the treatment and covariates reduces the imbalance by 79%.

Usually, reweighting estimators requires some overlap between the domain of treatment probability density of the treated and that of the controls, which is known as the common support condition. To our knowledge there are no concrete suggestions in this strand of the literature on how to impose a common support condition for continuous treatment. We follow Flores et al. (2012) to informally gauge the extent of overlap in the supports of different levels of the treatment. We divide these values into five quintiles and for each quintile, we compute the value of the GPS for each individual at the median level of the treatment for the quintile. We then compute the value of the GPS at the same median level of the treatment for the rest of individuals that are not part of that quintile. Finally, we compare the supports of the values of the GPS for these two groups (individuals in the quintile in question and the rest) by superimposing the densities. The plots referring to the five quintiles are not shown for space reasons. Fortunately, the evidence suggests that the support condition is likely to be satisfied in our sample. As a result this exercise does not imply any restriction on our sample (Flores et al., 2012), so that overall we are confident that reweighting eliminates any bias in the effects of the amount of CTs generated by a misspecification of the functional form of the outcome model.

These checks offer some evidence in favor of the validity of the assumptions of our benchmark model and hence of the reliability of its estimates. However, to consolidate our point that these estimates do not suffer from endogeneity bias we provide results from an instrumental variable approach where we instrument the household level treatment with two separate exclusion restrictions. The first stage regression estimates of these models are not reported to save on space, but they are available on request. Both variables are relevant instruments in all specifications as they are statistically significantly (at the one percent level) associated with household's i reported transfer. An F-test always rejects the null that all coefficients in the first stage are jointly zero.

Figure 10 shows the DRF curves for the whole sample, when the exclusion restriction used is the time needed to collect the transfer. In figure 11 the DRF curves are obtained by using the average transfer size at the community level as an instrument. We first compare the curves from the OLS approach in figure 4 to the corresponding curves from the IV approach in figures 10 and 11. The curves estimated by IV retain all the characteristics of the OLS curves, regardless of the instrument being used. The DRF curve for paid labor is the exact specular image of the DRF for on-farm labor which validates the labor reallocation interpretation. Compared to the OLS results, the IV DRFs for paid labor have a more pronounced decreasing pattern around the central range of treatment values where most of the observations are concentrated. As households decrease the hours they dedicate to paid labor at the same time they increase time dedicated to their own farms, as shown by the DRF on the right, and the intensity of effects increases with the amount of transfer received.[3]

Discussion and conclusions

A common perception among many policy makers, civil servants and more widely among the general public, is that cash transfers foster dependency of the vulnerable households who benefit from them. These political economy concerns have restricted either the adoption of cash transfer programs as a social protection instrument, or the scale of cash transfer provision, by both narrowing eligibility criteria and the value of the cash transfer (McCord, 2009). A transfer value between 10 to 30 percent of the median monthly household consumption expenditure has in fact been adopted in many programs in Africa (Daidone et al., 2017), irrespective of national or local income levels. This practice owes more to concerns about dependency than poverty reduction (see the example of Kenya in Pearson and Alviar, 2009). Further, these worries have only been partially supported by economic theory predictions, which in the traditional income-leisure approach posit a monotonically declining relationship between changes in non-labor income represented by the transfer and changes in supply of labor. However, as suggested by

[3] We do not show the DRF curves for hired labor for space reasons. They do not add much to the OLS analysis and are available from the authors upon request.

another branch of the literature, the relationship between labor supply and non-labor income can take the form of an inverted-U shape: positive up to some critical level of the transfer size, due to the higher utility of work brought about by increased consumption, negative after the identified threshold, because of the diminished utility due to leisure. From the perspective of the policy maker, it is therefore very relevant understanding, which is the level of the transfer inducing a potentially negative behavior from program beneficiaries.

In this study, we used data from a social cash transfer program implemented in Zambia, the Child Grant, and analyzed the effects of the amount of the cash transfer on three household-level labor outcomes: off-farm paid work, own farm work and agricultural labor hired in. We used recently developed econometric methods that allowed us to capture the elasticity of labor supply to non-labor income. Our findings show a shift from wage labor, which is mainly related to agricultural activities, to own farm labor. The story of switching from agricultural wage labor of last resort to on-farm activities emerged also from the qualitative work conducted in six countries in sub-Saran Africa (Barca et al., 2016). Agricultural wage labor in rural areas is often considered as a “refuge” sector where poor households work to hedge against agricultural risk or obtain needed liquidity. A reduction in participation and time worked in these activities can be deemed as welfare-enhancing.

We found that the Child Grant led to a reduction in the supply of beneficiary household labor to off-farm paid work regardless of the level of CT received by the household, though the effect is significant for relatively lower levels of the transfer. On the other hand, the cash transfer led to an increase in the labor supply to own farm labor activities, though significant only at relatively higher levels of transfers. The effect of the transfer on own farm labor activities becomes negative only above 1250 ZMK, suggesting the possible existence of an optimal transfer level (in terms of labor incentives) greater than the average transfer amount currently received by households enrolled in the Child Grant, and in any case suggesting that possible disincentive effects are well beyond the current transfer level (over 50 percent greater than the theoretical maximum). A similar effect is found on agricultural labor hired in by beneficiary households—it increases with the level of transfer, and is significant at higher levels of the transfer. At extremely high levels of transfers (far beyond the theoretical maximum) hired in labor appears to substitute for household labor supply to own farm activities.

The results of this study suggest that the Child Grant is unlikely to induce reduced work effort in targeted rural Zambian households, even at levels of transfer much higher than those currently being provided. Therefore, given the availability of a fiscal space or resource from donors, government policy in Zambia on social cash transfers, and more generally in similar development contexts, should avoid limiting the transfer size to low levels, which may actually undermine a significant impact on poverty reduction, which is the very purpose of a social protection program like a cash transfer.

References

Alzúa, M., Cruces, G., Ripani, L., 2013. Welfare programs and labor supply in developing countries: experimental evidence from Latin America. Journal of Population Economics 26(4), 1255–1284.

American Institutes for Research, 2013. 24-Month Impact Report for the Child Grant Programme. Washington DC: Author.

Ardington, C., Case, A., Hosegood, V., 2009. Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa. American Economic Journal: Applied Economics 1(1), 22–48.

Asfaw, S., Davis, B., Dewbre, J., Handa, S., Winters, P., 2014. Cash Transfer Programme, Productive Activities and Labor Supply: Evidence from a Randomised Experiment in Kenya. Journal of Development Studies 50(8), 1172–1196.

Barca, V., Brook, S., Holland, J., Otulana, M., Pozarny, P., 2015. Qualitative research and analyses of the economic impacts of cash transfer programmes in Sub-Saharan Africa. Synthesis Report. Food and Agriculture Organization of the United Nations: Rome.

Cerulli, G., 2014. Ctreatreg: Stata Module for Estimating Dose-Response Models under Exogenous and Endogenous Treatment. Working Paper Cnr-Ceris, N.05.

Daidone, S., Davis, B., Handa, S., Winters, P., 2017. The household and individual-level economic impacts of cash transfer programmes in sub-Saharan Africa. Synthesis Report. PtoP project report. Food and Agriculture Organization: Rome.

Del Carpio, X.V., 2008. Does Child Labor Always Decrease with Income? An Evaluation in the Context of a Development Program in Nicaragua, Policy Research Working Paper 4694. World Bank, Washington DC.

Fichera, E., Emsley, R., Sutton, M., 2016. Is treatment “intensity” associated with healthier lifestyle choices? An application of the dose response function. Economics and Human Biology 23, 149–163.

Fields, G., Hernández, R.D., Rodriguez, S.F., Sanchez Puerta, M.L., 2007. Earnings mobility in Argentina, Mexico, and Venezuela: Testing the divergence of earnings and the symmetry of mobility hypotheses. IZA Discussion Paper Series n. 3184. Bonn: IZA.

Fink, G., Jack, B.K., Masiye, F., 2014. Seasonal credit constraints and agricultural labor supply: Evidence from Zambia, NBER Working Paper 20218. National Bureau of Economic Research, Cambridge, MA.

Flores, C., Flores-Lagunes, A., Gonzalez, A., Neumann, T., 2012. Estimating the Effects of Length of Exposure to Instruction in a Training Program: The Case of Job Corps. Review of Economics and Statistics 94(1), 153–171.

Foguel, M., Barrios, R. Paes de, 2010. The Effects of Conditional Cash Transfer Programmes on Adult Labour Supply: An Empirical Analysis Using a Time-Series-Cross-Section Sample of Brazilian Municipalities. Estudos economicos 40(2), 259–293.

Fong, C., Imai, K., 2014. Covariate Balancing Propensity Score for General Treatment Regimes.

Gilligan, D.O., Hoddinot, J., Taffesse, A.S., 2009. The Impact of Ethiopia’s Productive Safety Net Program and its Linkages. Journal of Development Studies 45(10), 1684–1706.

Imai, K., Dyk, D.A. van, 2004. Causal Inference With General Treatment Regimes: Generalizing the Propensity Score. Journal of the American Statistical Association 99(467), 854–866.

Kang, J.D.Y., Schafer, J.L., 2007. Demystifying Double Robustness: A Comparison of Alternative Strategies for Estimating a Population Mean from Incomplete Data. Statistical Science 22(4), 523–539.

McCord, A., 2009 Cash Transfers and Political Economy in sub-Saharan Africa. ODI Project Briefing n. 31. Overseas Development Institute: London.

Pearson, R., Alviar, C., 2009. The Government of Kenya's cash transfer programme for vulnerable children: Political choice, policy choice, capacity to implement and targeting, from conception to adolescence 2002-2009. Final Draft, 14 June.

Prifti, E., Estruch, E., Daidone, S., Davis, B., Van Ufford, P., Michelo, S., Handa, S., Seidenfeld, D., Tembo, G., 2016 Learning about labour impacts of Cash Transfers in Zambia. Journal of African Economies *forthcoming*

Rawlings, L.B., Rubio, G.M., 2003. Evaluating the impact of conditional cash transfer programs. Lessons from Latin America. Policy research working paper 3119,The World Bank.

Ribas, R., Veras Soares, F., 2011. Is the effect of conditional transfers on labor supply negligible everywhere?

Robins, J.M.., Hernan, M.A., Brumback, B., 2000. Marginal structural models and causal inference in epidemiology. Epidemiology 11(5), 550–560.

Rose, E., 2001. Ex ante and ex post labor supply response to risk in a low-income area. Journal of Development Economics 64(2), 371–388.Skoufias, E., Di Maro, V., 2008. Conditional Cash Transfers, Adult Work Incentives, and Poverty. Journal of Development Studies 44(7), 935–960.

Uysal, D.S., 2015. Doubly Robust Estimation of Causal Effects with Multivalued Treatments: An Application to the Returns to Schooling. Journal of Applied Econometrics 30(5), 763–786.

Tables and Figures

Table 1: Characteristics of the sample at baseline by treatment status

	Original sample		
	Controls	Treated	Diff
HH members under 5y	1.905	1.881	0.003
	(0.746)	(0.773)	
HH members 6-12y	1.283	1.254	0.009
	(1.159)	(1.118)	
HH members 13-17y	0.438	0.494	-0.062*
	(0.719)	(0.75)	
Male HH members 18-64y	0.89	0.917	-0.051
	(0.647)	(0.618)	
Female HH members 18-64y	1.138	1.158	-0.02
	(0.454)	(0.529)	
Male HH members over 64y	0.017	0.018	0
	(0.13)	(0.134)	
HH size	5.69	5.746	-0.124
	(2.087)	(2.1)	
Number of orphans	0.311	0.297	0.017
	(0.959)	(0.971)	
Age of HH head	29.54	29.57	-0.349
	(9.081)	(9.128)	
Education of HH head	3.99	4.341	-0.318*
	(3.284)	(3.403)	
Dependency ratio	2.39	2.385	0.038
	(1.545)	(1.513)	
Price of maize	10.65	11.34	-0.621***
	(4.078)	(4.666)	
Price at potatoes	0.766	0.767	0.025
	(0.515)	(0.562)	
Number of cattle	0.363	0.593	-0.148
	(1.631)	(6.279)	
Number of poultry	2.053	2.144	-0.064
	(4.236)	(3.97)	
Operated land (ha)	0.479	0.522	-0.03
	(0.634)	(0.962)	
Female headed HH	0.996	0.994	0.005
	(0.062)	(0.078)	
Widowed HH head	0.055	0.052	0
	(0.228)	(0.222)	
Elderly HH head	0.007	0.005	0.001
	(0.082)	(0.071)	

Note: *** significant at 1%; ** significant a 5%, *significant at 10%.

Figure 1: Labor supply and cash transfers

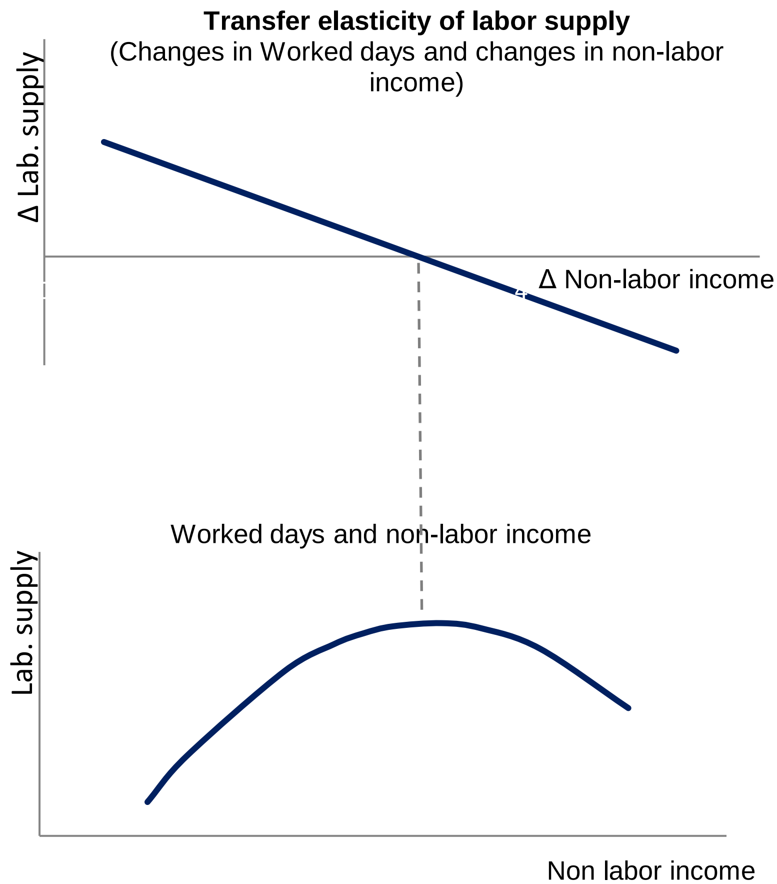

Figure 2: Probability density of the cash transfers

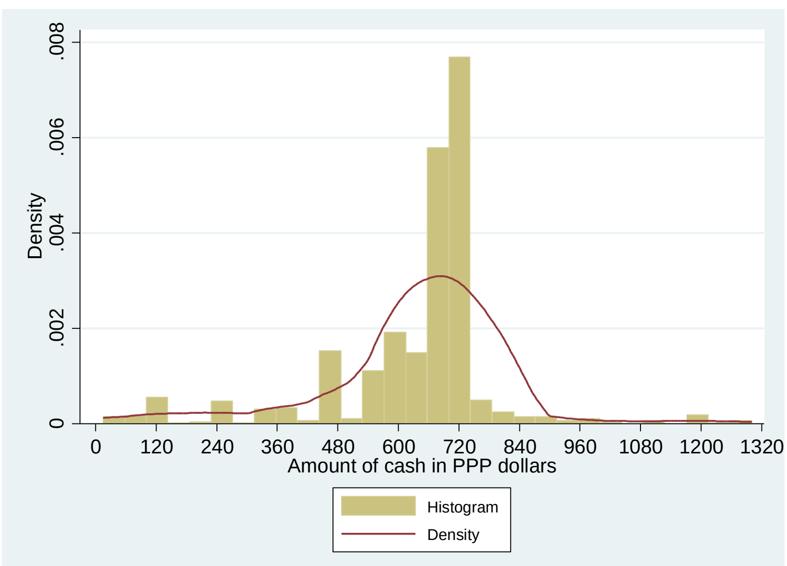

Figure 3: Average amount of supplied labor at different treatment levels

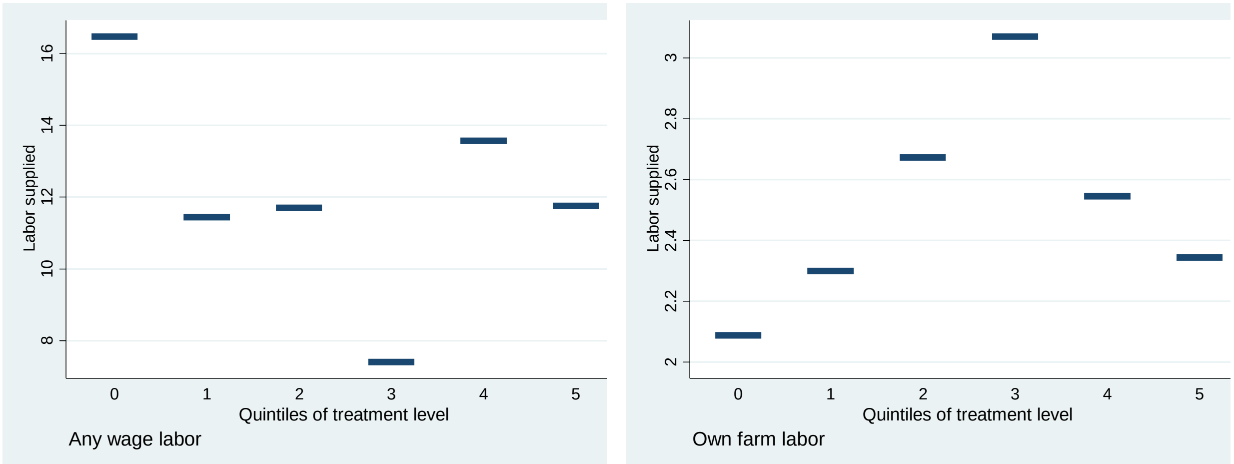

Figure 4: Response of labor supplied by HHs to changes in non-labor income

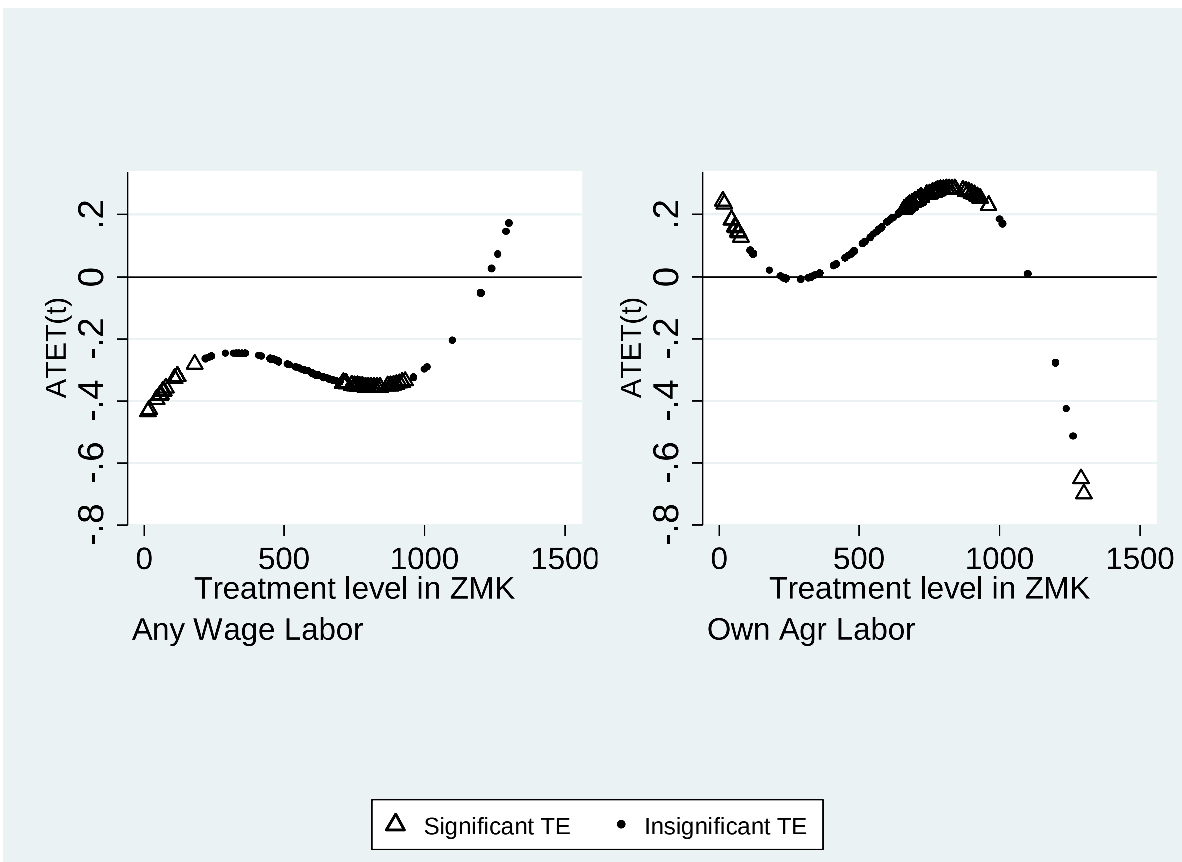

Figure 5: Relationship between supplied labor and non-labor income

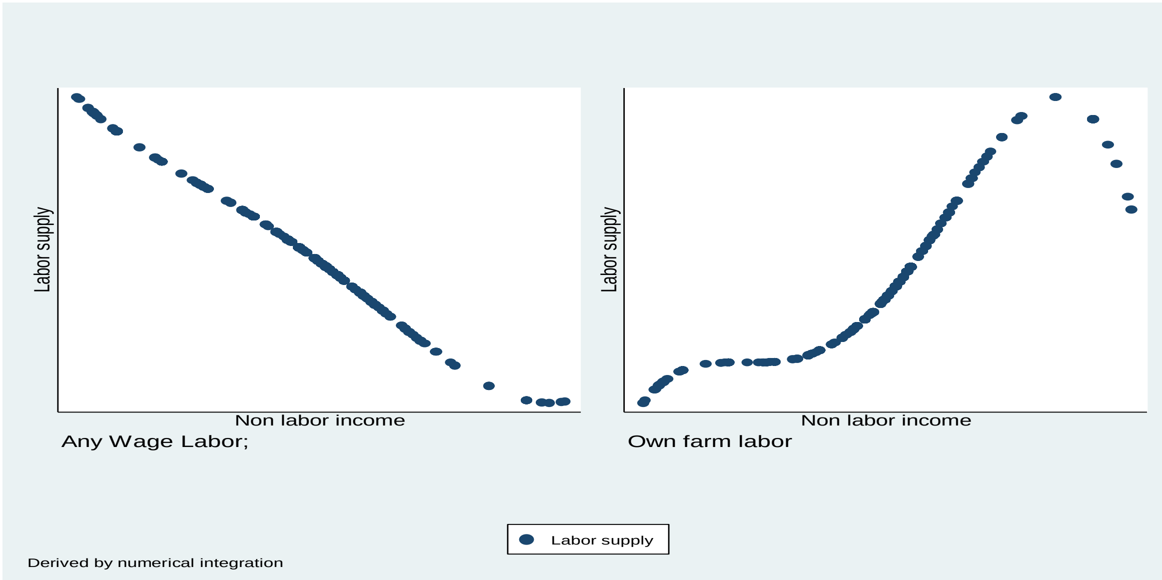

Figure 6: Standard normal quantile plots

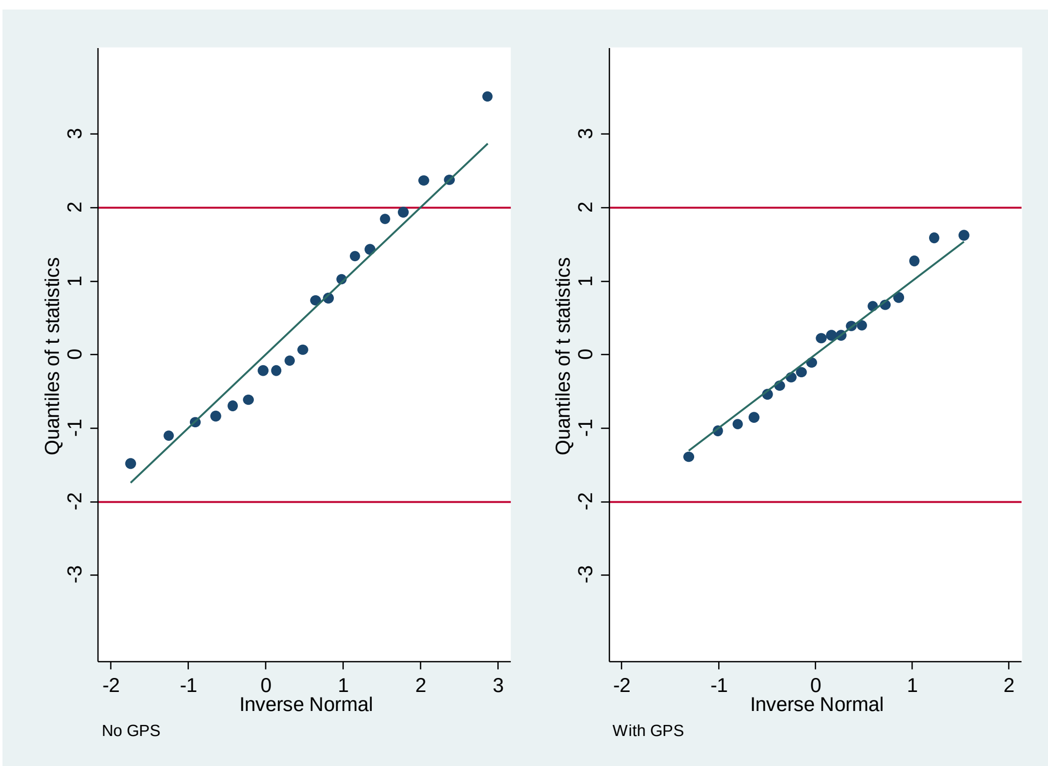

Figure 7: Response of labor supplied by HHs to changes in non-labor income: first IV estimates

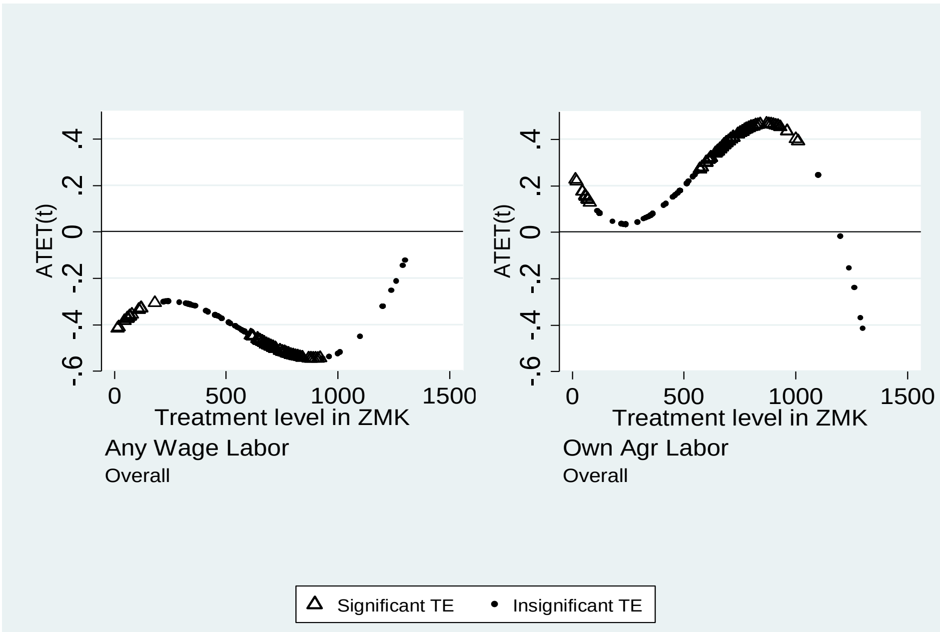

Figure 8: Response of labor supplied by HHs to changes in non-labor income: second IV estimates