

Household Response to Income Changes: Pre-Analysis Plan*

Johannes Haushofer[†], Jeremy Shapiro[‡]

November 21, 2015

Abstract

This document describes follow-up analyses to be performed on the data reported in Haushofer and Shapiro (2013), registered as <https://www.socialscisceregistry.org/trials/19>. The original paper describes a randomized controlled trial evaluating the Unconditional Cash Transfer (UCT) program of the NGO *GiveDirectly, Inc.* Between June 2011 and January 2013, *GiveDirectly* distributed unconditional cash transfers to 500 randomly selected poor rural households in Western Kenya. This document identifies methods for improving the robustness of the spillover analysis included in Haushofer and Shapiro (2013). It also includes several other minor additions.

JEL Codes: C93, D13, I15, I25, O12

Keywords: unconditional cash transfers, randomized controlled trial, impact evaluation.

*We are grateful to Jim Reisinger for excellent research assistance, and to Ingvild Almås, Mingyu Chen, Janet Currie, Jon de Quidt, Alexis Grigorieff, Larry Katz, Ilyana Kuziemko, David S. Lee, Chris Roth, Dan Sacks, Simone Schaner, and Tom Vogl for comments and discussion. This research was supported by NIH Grant R01AG039297 and Cogito Foundation Grant R-116/10 to Johannes Haushofer.

[†]Peretsman Scully Hall 427, Princeton University, Princeton, NJ 08544, and Busara Center for Behavioral Economics, Nairobi, Kenya. haushofer@princeton.edu

[‡]Busara Center for Behavioral Economics, Nairobi, Kenya. jeremy.shapiro@busaracenter.org

1 Introduction

The present document introduces additional analyses on the dataset first reported in Haushofer and Shapiro (2013). Our primary focus is on outlining methods to improve the robustness of the spillover analysis in that paper, and addressing the fact that households in control villages were selected for the participation close to one year after the households in treatment villages were selected.

2 Analysis and Econometric Specifications

2.1 Adjusting for thatched roof selection criterion

A potential weakness in the spillover analysis reported in Haushofer and Shapiro (2013) is that the thatched-roof selection criterion for participation in the study was applied to households in control villages one year after it was applied to households in treatment villages. As a result, there is endogenous selection into the pure control condition, as some proportion of households in pure control villages are likely to have upgraded to a metal roof over this time period, and these households are excluded from endline in the pure control villages. These households are potentially different both from households that did not upgrade, and different from households in treatment villages that only upgraded in response to their neighbors receiving transfers. Thus, the fact that no metal roof households are included in the endline survey of control villages potentially biases the spillover analysis. In the following, we describe the selection problem formally, and outline the analyses we will perform to bound any resulting bias.

2.1.1 Basic Selection Problem

We treat the bias introduced by the time lag in the application of the thatched roof criterion as selection bias. Consider the following sample selection model (cf. Angrist, Bettinger, and Kremer 2006 and Lee 2009):

D = Assignment to treatment village

S = Household takes the endline survey

S_1 = Household would take the endline survey if assigned to spillover status

S_0 = Household would take the endline survey if assigned to pure control status

Y = Outcome of interest

Y_1 = Outcome of interest if assigned to spillover status

Y_0 = Outcome of interest if assigned to pure control status

Note that the sample is restricted to *control* households (in both treatment and pure control villages), and the treatment dummy D identifies *spillover* households, i.e. control households in treatment villages. For now we abstract away from selection bias through attrition and consider only bias from differential application of the thatched roof eligibility criterion. For each individual, we only observe one of the sample selection indicators S_1 , S_0 . Similarly, Y_0 and Y_1 are latent potential outcomes that we only observe if an individual takes the endline survey. Thus:

$$S = S_1D + S_0(1 - D)$$

$$Y = S [Y_1D + Y_0(1 - D)]$$

Randomization gives us $Y_0, Y_1, S_0, S_1 \perp D$. Calculating the spillover effect from the observed sample gives us $E[Y | S = 1, D = 1] - E[Y | S = 1, D = 0]$. However, this is a biased measure of the average effect of living in a treatment village for individuals who were observed:

$$\begin{aligned} & E[Y | S = 1, D = 1] - E[Y | S = 1, D = 0] \\ &= E[Y_1 | S_1 = 1, D = 1] - E[Y_0 | S_0 = 1, D = 0] \\ &= E[Y_1 | S_1 = 1] - E[Y_0 | S_0 = 1] \\ &= E[Y_1 | S_1 = 1] + E[Y_1 | S_0 = 1] - E[Y_1 | S_0 = 1] - E[Y_0 | S_0 = 1] \\ &= E[Y_1 - Y_0 | S_0 = 1] + E[Y_1 | S_1 = 1] - E[Y_1 | S_0 = 1] \\ &= E[Y_1 - Y_0 | S_0 = 1] + E[Y_1] - E[Y_1 | S_0 = 1] \end{aligned}$$

The term $E[Y_1] - E[Y_1 | S_0 = 1]$ identifies any bias arising from the fact that individuals who did not upgrade their roofs (and thus were eligible to be surveyed) in pure control villages may have different outcomes from those who do. Also note that we use $E[Y_1 | S_1 = 1] = E[Y_1]$, since all households in treatment villages are observed. We will perform the following analyses to bound this selection effect.

2.1.2 Spillover effect including metal roof households

Identifying assumption 1: Random selection into roof upgrade We begin by assuming that selection into roof upgrade (and hence out of the endline survey in the pure control villages) is random, i.e. $Y \perp S_0$, and therefore

$$E[Y_1] = E[Y_1 | S_1 = 1] = E[Y_1 | S_0 = 1].$$

This assumption allows us to identify the treatment effect through the simple comparison of all households that took the endline survey, i.e. $E[Y_1 - Y_0 | S_0 = 1] = E[Y | S = 1, D = 1] - E[Y | S = 1, D = 0]$. We refer to this as the “naïve” analysis.

To provide evidence in support of this assumption, we would ideally ask whether selection into upgrade can be predicted from baseline observables in the pure control group; however, we do not have data on the metal roof households. A second-best option is to ask whether selection into upgrade can be predicted from baseline observables in the spillover group. We will do this using the following specification:

$$U_{hv} = \beta_0 + \beta_1 X_{\{i\}hvB} + \varepsilon_{\{i\}hv}$$

Here, $X_{\{i\}hvB}$ is a vector of baseline characteristics of respondent i (if measured at the individual level) in household h in village v at baseline ($t = B$). U_{hv} is an indicator variable taking the value of 1 if household h upgraded to a metal roof between baseline and endline and 0 otherwise. Note that we will exclude treatment and pure control households from this analysis. $\varepsilon_{\{i\}hv}$ is an idiosyncratic error term. Standard errors are clustered at the household level. β_1 identifies the extent to which baseline characteristics predict upgrade to metal roof, and thus whether selection into upgrade can be considered random with respect to outcome variables. We will use the eight index variables as predictors of upgrade.

A final source of evidence for the comparability of the spillover and pure control samples is to compare them on baseline characteristics. However, no baseline survey was administered to pure control households. Nevertheless, there some individual and household characteristics are either immutable or calculable from endline values. We will determine whether these characteristics are balanced between spillover and pure control households using the following

specification:

$$y_{\{i\}hvB} = \beta_0 + \beta_1 Sp_{hv} + \varepsilon_{\{i\}hvB}$$

Here, $y_{\{i\}hvB}$ is a characteristic of respondent i (if measured at the individual level) in household h in village v at baseline ($t = B$). Sp_{hv} is an indicator variable taking the value of 1 if household h is a spillover household and 0 if it is a pure control household. Note that we will exclude treatment households from this analysis. $\varepsilon_{\{i\}hvB}$ is an idiosyncratic error term. Standard errors are clustered at the village level. β_1 identifies differences in immutable characteristics between spillover and pure control households.

We will use the following characteristics as comparison variables:

1. Age of primary respondent
2. Gender of primary respondent
3. Marital status of primary respondent at baseline
4. Highest level of education attained by primary respondent
5. Number of children, excluding those born between baseline and endline

In a separate analysis, we will ascertain that these characteristics are truly immutable by comparing baseline and endline among spillover households.

2.1.3 Controlling for baseline characteristics

Identifying assumption 2: Random selection into roof upgrade conditional on observables We next assume that selection of spillover and pure control households into upgrade (and hence the endline survey) is random conditional on a set of observable household characteristics X . In this case, the conditional independence assumption holds, i.e. $Y \perp S_0 \mid X$. Thus, if we control for these covariates in our specification, we can identify the treatment effect:

$$\begin{aligned} & E[Y_1 - Y_0 \mid S_0 = 1, X] + E[Y_1 \mid S_1 = 1, X] - E[Y_1 \mid S_0 = 1, X] \\ &= E[Y_1 - Y_0 \mid S_0 = 1, X] + E[Y_1 \mid X] - E[Y_1 \mid X] \\ &= E[Y_1 - Y_0 \mid S_0 = 1, X] \\ &= E[Y_1 - Y_0 \mid S_0 = 1] \end{aligned}$$

The last equality is true because $Y \perp S_0 \mid X$ by assumption. Thus, we will recalculate the spillover effect using the baseline characteristics listed above as control variables. We will use the following specification:

$$y_{\{i\}hvE} = \beta_0 + \beta_1 S_{hv} + X_{\{i\}hv} \gamma + \varepsilon_{\{i\}hvE}$$

Here, $y_{\{i\}hvE}$ is an outcome of interest (using those outcomes reported in Haushofer and Shapiro 2013) for respondent i (if measured at the individual level) in household h in village v at baseline ($t = E$). S_{hv} is an indicator variable taking the value of 1 if household h is a spillover household and 0 if it is a pure control household. Note that we will exclude treatment households from this analysis. $X_{\{i\}hv}$ is a vector of individual and household level demographic variables listed above. $\varepsilon_{\{i\}hvE}$ is an idiosyncratic error term. Standard errors are clustered at the village level. β_1 identifies the spillover effect.

2.1.4 Restricting the sample to households with thatched roofs at endline

We now consider improvements in identification resulting from restricting the sample to households which still have thatched roofs at endline. To begin, we define notation for never-takers, always-takers, compliers, and defiers of metal roof upgrade between baseline and endline. Note that again the sample is restricted to households in the spillover and pure control groups, and therefore a complier household is one that upgrades to a metal roof as the result of a spillover effect from neighboring households receiving a transfer; a defier household is one that does not upgrade for this reason. We denote actual and potential roof status at endline as follows:

R = Household upgrades to metal roof between baseline and endline

R_1 = Household would upgrade to metal roof if assigned to spillover status

R_0 = Household would upgrade to metal roof if assigned to pure control status

We further denote the proportions of always-takers, compliers, defiers, and never-takers as follows:

$$\text{Always-takers: } Pr(R_0 = 1, R_1 = 1) = \alpha$$

$$\text{Compliers: } Pr(R_0 = 0, R_1 = 1) = \gamma$$

$$\text{Defiers: } Pr(R_0 = 1, R_1 = 0) = \phi$$

$$\text{Never-takers: } Pr(R_0 = 0, R_1 = 0) = \nu$$

Now consider the comparison of spillover and pure control households which still have

thatched roofs at endline. In treatment villages, the households with thatched roofs at endline are either defiers or never-takers. In pure control villages, they are either compliers or never-takers. The difference between the observed and potential outcomes can therefore be written and rearranged as follows:

$$\begin{aligned}
& \nu E[Y_N|S = 1, D = 1] + \phi E[Y_F|S = 1, D = 1] - \nu E[Y_N|S = 1, D = 0] - \gamma E[Y_C|S = 1, D = 0] \\
&= \nu E[Y_{N,1}|S_1 = 1, D = 1] + \phi E[Y_{F,1}|S_1 = 1, D = 1] - \nu E[Y_{N,0}|S_0 = 1, D = 0] - \gamma E[Y_{C,0}|S_0 = 1, D = 0] \\
&= \nu E[Y_{N,1}|S_1 = 1] + \phi E[Y_{F,1}|S_1 = 1] - \nu E[Y_{N,0}|S_0 = 1] - \gamma E[Y_{C,0}|S_0 = 1] \\
&= \nu E[Y_{N,1}|S_1 = 1] + \phi E[Y_{F,1}|S_1 = 1] - \nu E[Y_{N,0}|S_0 = 1] - \gamma E[Y_{C,0}|S_0 = 1] \\
&\quad + \nu E[Y_{N,1}|S_0 = 1] - \nu E[Y_{N,1}|S_0 = 1] \\
&= \nu E[Y_{N,1} - Y_{N,0}|S_0 = 1] + \phi E[Y_{F,1}|S_1 = 1] - \gamma E[Y_{C,1}|S_0 = 1]
\end{aligned}$$

Thus, the difference between households with thatched roofs at endline is identified for never-takers, except for the difference between the proportion and potential outcomes of households that are compliers or defiers in terms of upgrading to metal roofs. We next outline under which assumptions this bias is zero or can be bounded.

Identifying assumption 3: Monotonicity (“no defiers”) We first make the classic monotonicity or “no defiers” assumption that is at the foundation of many randomized field experiments (Angrist, Imbens, and Rubin 1996). In our framework, the assumption states that $\phi = 0$. Could there be defiers in our sample? In our view, the only plausible reason for control households to refrain from upgrading their thatched roofs to metal is to remain eligible for possible future transfers from *GiveDirectly*. However, control households in treatment villages were credibly told by GiveDirectly that they would not receive cash transfers. The no-defier assumption is therefore reasonable in our setting.

With this assumption, the only bias arises from compliers, which are included in the pure control thatched-roof sample but not in the spillover thatched-roof sample because they upgraded to metal roofs. Importantly, can find out how many such households there are by obtaining a precise estimate of the magnitude of the spillover effect of the cash transfers on metal roof ownership. In September 2015, we returned to households with metal roofs in pure control villages to ascertain when they upgraded to a metal roof. Households that upgraded between April 2011 and June 2012 should originally have been eligible for participation in the study, but were excluded because of the late application of the thatched roof criterion. We identified 170 such households. We then used the same algorithm originally used to select pure control households to calculate the probability that each of these households would

have been included in the study had they been identified as eligible at the time. The original sampling method required us to select 8 households from the pool of eligible households in each village (those with thatched roofs). When there were 8 or fewer eligible households in a given village, we selected all households. When there more than 8 eligible households, we selected 8 with equal probability for each. We were thus able to calculate the exact probability that a given household would be selected in each village. In villages with 8 or fewer eligible households, the probability of selection was 1. In villages with more than 8 eligible households, the probability was 8 divided by the total number of eligible households. To determine how many of the 170 “recall” households should have been selected for the survey, we multiply this probability by the number of recall households in each village, resulting in a total of 78 households. Since there were 432 pure control households in the original study, this gives us an upgrade rate from baseline to endline of $78/(432 + 78) = 0.153$ for pure control villages. Similarly, since there were a total of 469 spillover households at endline, of which 77 had metal roofs, the upgrade rate among spillover households was $77/469 = 0.164$. Applying the upgrade rate of 0.153 in pure control village to these spillover households, we would predict $0.153 \cdot 469 = 72$ metal roofs in the spillover group at endline. In actuality we observe 77 metal roof households. The treatment therefore had a spillover effect on metal roof ownership of $77 - 72 = 5$ households.

We take two approaches to the bias arising from these five households. The first is to ignore it: with 5 households out of 469, i.e. 1.1 percent, the spillover effect of transfers on metal roof ownership is negligible. We can therefore consider the spillover analysis that restricts the sample to households that still have thatched roof at endline as nearly uncontaminated by spillover effects on metal roof ownership. In this case, restricting the sample to households that still have thatch roofs at endline identifies the spillover effect. The second approach is to bound the spillover effect using worst-case assumptions. We will therefore report Lee and Manski bounds.

Identifying assumption 4A: Same proportion and potential outcomes for compliers and defiers We now relax the monotonicity assumption and ask under which alternative assumptions the comparison of thatch-at-endline households in treatment and control villages identifies the spillover effect. One such assumption is that the proportion and potential outcomes of compliers and defiers are the same, i.e.

$$\phi E[Y_{F,1}|S_1 = 1] = \gamma E[Y_{C,1}|S_0 = 1].$$

This assumption says that the proportion and outcome distribution of households which are induced to upgrade to metal roofs when their neighbors receive transfers are identical to those of households which are induced to keep their thatched roofs by treatment. That the outcome distribution of these two types is similar is plausible because both types are marginal, i.e. they are “ready to upgrade” before transfers.

Identifying assumption 4B: Same potential outcomes for compliers and $\frac{\gamma}{\phi}$ of the defiers A weaker assumption is that only $\frac{\gamma}{\phi}$ of the defiers have the same potential outcomes as the compliers. This leaves a proportion of $\phi - \gamma$ of the sample whose outcome distribution we don’t know and who therefore contaminate the spillover effect estimate. However, from the exercise described above, we know that $\phi - \gamma = 0.011$. Again, this is negligible and can either be ignored, or bounded as described above. The details of this approach have been described by Angrist, Imbens, and Rubin (1996) and de Chaisemartin (2013).

Identifying assumption 4C: Same potential outcomes for compliers, never-takers, and defiers Finally, we can relax the assumption that the proportion of compliers and defiers are the same if we instead assume that their distribution of potential outcomes is the same as that of the never-takers, i.e. $E[Y_{F,1}|S_1 = 1] = E[Y_{C,1}|S_0 = 1] = E[Y_{N,1}|S_0 = 1]$. This assumption says that the spillover group, which consists of never-takers and defiers, has the same outcome distribution as the pure control group, which consists of never-takers and defiers.

2.1.5 Testing whether inclusion vs. exclusion of metal roof households affects results

We next ask whether including vs. excluding households with metal roofs at endline from the spillover analysis affects results. To this end, we will analyze the difference in spillover effects when calculated across all spillover households and when excluding spillover households that upgraded. If we find that the results are similar whether or not we exclude metal roof households, this suggests that the differential application of the thatched roof criterion introduced only minimal bias. We will estimate a series of models of the form:

$$y_{\{i\}hvE,m} = \beta_m S p_{hv} + \varepsilon_{\{i\}hvE,m}$$

where m denotes the model number, $T_{hv,m}$ is an indicator variable for whether household h in village v of model m upgraded, and $\varepsilon_{\{i\}hvE,m}$ is an idiosyncratic error term. Note that h (or i for individual measures) indexes either the total number of spillover households or the

number of spillover households that did not upgrade:

1. If $m = 1$, then $h = 1 \dots H_1$ where H_1 is the total number of spillover households and pure control households.
2. If $m = 2$, then $h = 1 \dots H_2$ where H_2 is the total number of spillover households that did not upgrade and pure control households

Writing each specification in vector form and stacking, we get the seemingly unrelated regression model:

$$\begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} = \begin{pmatrix} Sp_1 & 0 \\ 0 & Sp_2 \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix} + \begin{pmatrix} \varepsilon_1 \\ \varepsilon_2 \end{pmatrix}$$

where Y_m is the vector of each $y_{\{i\}hvE,m}$ in model m , S_m is the vector of each Sp_{hv} in model m . ε_m is the vector of each $\varepsilon_{\{i\}hvE,m}$ in vector m . Standard errors are clustered at the village level.

Thus, β_m will identify the spillover effect in model m . After estimating this specification, we can test for the equality of each β_1 and β_2 . If we cannot reject $H_0 : \beta_1 = \beta_2$, this again suggests that spillover households that upgraded are not significantly different in terms of the outcome variable to those that did not upgrade, and that the late application of the exclusion restriction introduced minimal bias into our calculation of the spillover effect.

2.2 Re-analysis of within-village treatment effects with controls

We will also re-analyze within-village treatment effects in the main analyses in Haushofer and Shapiro (2013) using demographic and baseline measures as controls. For within-village treatment effects, we estimate the following model:

$$y_{\{i\}hvE} = +\beta_0 + \beta_1 T_{hv} + \delta_1 y_{\{i\}hvB} + \delta_2 M_{\{i\}hvB} + X_{\{i\}hv} \gamma + \alpha_v + \varepsilon_{\{i\}hvE} \quad (1)$$

Here, $y_{\{i\}hv}$ is the outcome of interest for household h in village v , measured at endline, of individual i . The sample is restricted to treatment and spillover households. Village-level fixed effects are captured by α_v . T_{vh} is a treatment indicator that takes value 1 for treatment households, and 0 otherwise. $X_{\{i\}hv}$ is a vector of individual and household level demographic variables listed below. Following McKenzie (2012), we condition on the baseline level of the outcome variable when available, $y_{\{i\}hvB}$, to improve statistical power. To include observations where the baseline outcome is missing, we code missing values as zero and include a dummy indicator that the variable is missing ($M_{\{i\}hvB}$). $\varepsilon_{\{i\}hvE}$ is an idiosyncratic error term.

Thus, β_1 identifies the treatment effect for treated households relative to spillover households. Standard errors are clustered at the level of the unit of randomization, i.e. the household. In addition to this standard inference, we compute FWER-corrected p -values across the set of index variables. Finally, we estimate the system of equations jointly using seemingly unrelated regression (SUR), which allows us to perform Wald tests of joint significance of the treatment coefficient across outcome variables.

The control variables are the following:

1. Demographic

- (a) Age of primary respondent at baseline
- (b) Gender of primary respondent at baseline
- (c) Marital status / household type (single vs. married) at baseline
- (d) Highest level of education attained by primary respondent at baseline
- (e) Number of children at baseline

2. Economic

- (a) Baseline consumption, asset levels, and land holdings at baseline
- (b) Ownership of non-agricultural enterprise at baseline
- (c) Ownership of agricultural enterprise at baseline
- (d) Participation in wage labor at baseline

2.3 Transfer and survey timing

To determine whether transfer timing had an impact on outcomes in the comparison of large transfer recipient households to small transfer recipient households, we will re-estimate the analysis that distinguishes recipients of large transfers from recipients of small transfers while controlling for the number of months between receipt of half of the transfers and endline:

$$y_{\{i\}hvE} = \alpha_v + \beta_0 + \beta_1 T_{hv}^L + \beta_2 T_{hv}^S + \delta_1 y_{\{i\}hvB} + \delta_2 M_{\{i\}hvB} + \tau_{hv} + \varepsilon_{\{i\}hvE} \quad (2)$$

Here, T_{hv}^x are indicator variables for either receipt of small transfers ($x = S$) or large transfers ($x = L$) by household h in village v . The sample is restricted to treatment villages. τ_{hv} is the number of months between the date at which half of the individual transfers to a household had been made and the endline.

In addition, to assess whether households which completed the survey early differed from those which completed it late, we will compare the endline date to baseline village and household characteristics, separately for treatment, spillover, and pure control households (restricting to fixed endline characteristics for the latter group):

$$SM_{\{i\}hv}^E = \beta_0 + X_{\{i\}\{h\}vB}\beta + \varepsilon_{\{i\}hv}$$

$SM_{\{i\}hv}^E$ is the number of months from the date that the first endline survey was conducted to the date for which the endline survey was administered to individual i (omitted for measures from the household survey) in household h in village v . $X_{\{i\}\{h\}vB}$ is a vector of baseline household and individual characteristics for individual i in household h of village v . $\varepsilon_{\{i\}hv}$ is an idiosyncratic error term. Thus the vector β captures any correlation between baseline characteristics and the timing of the endline survey.

In addition, we will ask whether treatment status predicts when a household completed endline:

$$SM_{\{i\}hv}^E = \beta_0 + \beta_1 T_{hv} + \beta_2 Sp_{hv} + \varepsilon_{\{i\}hv}$$

Standard errors will be clustered at the village level. β_1 captures the difference between treatment and pure control households in survey month date. β_2 captures the difference in endline month between spillover and pure control households.

2.4 Labor outcomes

The original paper contained a non-pre-specified dummy variable for wage labor being the primary source of income. To get a more fine-grained measure of the effects of the program on labor supply, we will additionally report treatment effects on the following variables:

- Dummy for salaried job being the primary source of income
- Proportion of working-age household members who spent any time in the last 12 months doing casual labor
- Proportion of working-age household members who spent any time in the last 12 months doing a salaried job
- Number of income-generating activities
- Amount spent on hiring labor for agricultural activities

2.5 Political Outcomes

To assess any treatment effects on political engagement, we will analyze the following outcome variables, and an index that combines all of them:

- Will you be voting in the upcoming national elections that will be held next year?
- Do you know the names of all the candidates who will be running for Prime Minister and President in next year’s elections?
- It is every Kenyan citizen’s responsibility to vote
- Kenya receives 20 billion shillings per year for development from foreigners, what do you think is the best way to decide how to use that money to reduce poverty? Dummy for responding “Let the Kenyan government decide how to spend it”

2.6 Re-analysis of effects by gender of respondent

We will reanalyze treatment effects by the gender of the *respondent* (rather than recipient), restricting the sample to cohabiting households only. Note that this analysis was already conducted during data analysis for the first paper, but was not included in the first draft of the working paper. It is included here to signal a commitment to make the results available. We emphasize that these analyses were *not* pre-specified before being first conducted.

2.7 Distinguishing investment and durable investment

Finally, we will distinguish between investment in durables and non-durables by generating variables for total spending on such durables, as follows:

Durable investment:

1. Livestock (cows / bulls, sheep, goats, pigs, birds)
2. Machinery and durable goods for enterprises
3. Farm implements (tools, wheelbarrows, cars, etc.)
4. Home improvement (roof, building materials, pit latrine etc.)
5. Transportation (motobikes, bicycles)

Non-durable investment:

1. Agricultural inputs (seed, fertilizer, water, hired labor, livestock feed, livestock medicine etc.)
2. Enterprise expenses (wages, electricity, water, transport, inventory other inputs)
3. Education (school and college fees, books, uniforms)
4. Savings

References

- Angrist, Joshua, Eric Bettinger, and Michael Kremer. 2006. “Long-Term Educational Consequences of Secondary School Vouchers: Evidence from Administrative Records in Colombia.” *The American Economic Review* 96 (3): 847–862.
- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. “Identification of Causal Effects Using Instrumental Variables.” *Journal of the American Statistical Association* 91 (434): 444–455.
- de Chaisemartin, Clement. 2013. “Defying the LATE? Identification of local treatment effects when the instrument violates monotonicity.” The Warwick Economics Research Paper Series (TWERPS) 1020, University of Warwick, Department of Economics.
- Haushofer, Johannes, and Jeremy Shapiro. 2013. “Household Response to Income Changes: Evidence from an Unconditional Cash Transfer Program in Kenya.” *Working Paper*.
- Lee, David S. 2009. “Training, wages, and sample selection: Estimating sharp bounds on treatment effects.” *The Review of Economic Studies* 76 (3): 1071–1102.
- McKenzie, David. 2012. “Beyond baseline and follow-up: The case for more T in experiments.” *Journal of Development Economics* 99 (2): 210–221.