General Equilibrium Effects of Cash Transfers: Pre-Analysis Plan for Household Spillovers

Johannes Haushofer,* Edward Miguel,† Paul Niehaus,‡ and Michael Walker§

June 20, 2018

1 Introduction

This document is an amendment to the Household welfare PAP filed for the General Equilibrium Effects of Cash Transfers in Kenya (GE) project [Haushofer et al., 2017, filed July 6, 2017]. In that analysis plan, we had specified only specifications analyzing direct treatment effects on treated households. This document additional discusses the analysis of treatment effects on non-recipient and ineligible households, and the effects of saturation on both recipient, non-recipient, and ineligible households. For details on the intervention and experimental design, please see the household welfare PAP [Haushofer et al., 2017].

2 Data and variable construction

The data source for our investigation of household spillovers comes from household survey data. In advance of the distribution of transfers to a treatment village, we conducted a baseline household census and household survey. The household census was designed to be comparable to the census conducted by the partner organization GiveDirectly (GD), but to ensure there was no systematic bias between their censusing methods and ours, we conducted our own censuses in all villages (both treatment and control). The census served as a sampling frame for baseline household surveys. We determined household eligibility based on the census data and randomly selected 8 eligible and 4 ineligible households to

---
*Princeton University, NBER, and Busara Center for Behavioral Economics, haushofer@princeton.edu
†UC Berkeley, emiguel@berkeley.edu
‡UC San Diego, pniehaus@ucsd.edu
§UC Berkeley, mwwalker@berkeley.edu

be targeted for surveys (12 total) in each village; we refer to these households as "initially-sampled" households. For couples, we randomly selected either the male or female to be the "target" respondent; if we could not reach the target, but the spouse/partner was available, we surveyed the spouse/partner. If an initially-sampled household was not available to be surveyed on the day we visited the village for baseline surveys, we replaced this household with another randomly-selected household to ensure that we surveyed 12 households in each village. We refer to these households as "replacement" households. Lastly, we refer to households that were initially-sampled but unable to be surveyed as "missed baseline" households.

Endline surveys target all "initially sampled" and "replacement" households. For households that were baselined, we attempt to survey the same respondent that was surveyed at baseline. Endline surveys began at the end of May 2016 and ran through May 2017. The median survey date was about 1.5 years after the baseline surveys. See Haushofer et al. [2017] for more details on endline data collection.

3 Empirical Specifications

3.1 Saturation effects

To estimate the effect of saturation, we use the following specification:

$$
\begin{aligned}
y_{ihvs,t=1} = {} & \beta_0 + \beta_1 T_{vs} + \beta_2 E_{hvs} + \beta_3 H_s + \beta_4 T_{vs} \times E_{hvs} + \beta_5 T_{vs} \times H_s \\
& + \beta_6 E_{hvs} \times H_s + \beta_7 T_{vs} \times E_{hvs} \times H_s + \delta_1 y_{ihvs,t=0} + \delta_2 M_{ihvs} + \varepsilon_{ihvs}
\end{aligned}
\tag{1}
$$

Here, h indexes the household, v indexes the village, s indexes the sublocation, and t indicates whether the variable was measured at baseline or endline. For individual-level outcomes, i indexes household members. T_{vs} is an indicator for households residing in a treated village, E_{hvs} is an indicator for whether the household is eligible for transfers, and H_s is an indicator for living in a high-saturation sublocation. Following McKenzie [2012], we condition on the baseline values of the outcome variable $y_{hv,t=0}$ to improve statistical power. When $y_{hv,t=0}$ is missing for an observation, we include an indicator term for missingness M_{ihvs} and replace $y_{hv,t=0}$ with its mean. We cluster standard errors at the saturation group level to estimate the effects of saturation. Inverse probability weights will be used to account for the relative frequency of eligible and ineligible households and slight

deviations from the target number of households in each village.

This specification allows us to estimate the following effects, clustering standard errors at the saturation group level:

1. Treated eligible saturation regression: Effect of saturation on eligible households in treatment villages:
$E[y_{ihvs,t=1} \mid T = 1, E = 1, H = 1] - E[y_{ihvs,t=1} \mid T = 1, E = 1, H = 0] = \beta_3 + \beta_5 + \beta_6 + \beta_7$

2. Untreated eligible saturation regression: Effect of saturation on eligible households in control villages:
$E[y_{ihvs,t=1} \mid T = 0, E = 1, H = 1] - E[y_{ihvs,t=1} \mid T = 0, E = 1, H = 0] = \beta_3 + \beta_6$

3. Treated ineligbile saturation regression: Effect of saturation on ineligible households in treatment villages
$E[y_{ihvs,t=1} \mid T = 1, E = 0, H = 1] - E[y_{ihvs,t=1} \mid T = 1, E = 0, H = 0] = \beta_3 + \beta_5$

4. Untreated ineligbile saturation regression: Effect of saturation on ineligible households in control villages
$E[y_{ihvs,t=1} \mid T = 0, E = 0, H = 1] - E[y_{ihvs,t=1} \mid T = 0, E = 0, H = 0] = \beta_3$

5. Pooled eligible saturation regression: Effect of saturation on eligible households in treatment and control villages: Average of 1. and 2., weighted by share of treatment vs. control villages (~50%).

6. Pooled ineligible saturation regression: Effect of saturation on ineligible households in treatment and control villages: Average of 3. and 4., weighted by share of treatment vs. control villages (~50%).

7. Pooled treated saturation regression: Effect of saturation on eligible and ineligible households in treatment villages: Average of 1. and 3., weighted by share of eligible vs. ineligible households.

8. Pooled untreated saturation regression: Effect of saturation on eligible and ineligible households in control villages: Average of 2. and 4., weighted by share of eligible vs. ineligible households.

9. Pooled saturation regression: Effect of saturation on eligible and ineligible households in treatment and control villages: Average of 1.–4., weighted by share of treatment vs. control villages and eligible vs. ineligible households. This is our main analysis of interest for saturation.

3.2 Spillover effects

To estimate spillover effects, we use the following specification, clustering standard errors at the village level:

$$y_{ihvs,t=1} = \beta_0 + \beta_1 T_{vs} + \beta_2 H_s + \delta_1 y_{ihvs,t=0} + \delta_2 M_{ihvs} + \varepsilon_{ihvs} \qquad (2)$$

Note this is the same specification as that used to study treatment effects in the household PAP. The difference is that to identify spillover effects, the sample will be restricted to ineligible households. In addition, the specification will be estimated both for high-saturation sublocations alone and low-saturation sublocations alone (omitting the H_s dummy); and for all ineligible households in both high-saturation and low-saturation sublocations, weighting appropriately according to the relative frequency of households.

This specification allows us to estimate the following effects:

1. Ineligible spillover regression, high saturation: Effect on ineligibles of residing in a treatment village in a high-saturation sublocation.

2. Ineligible spillover regression, low saturation: Effect on ineligibles of residing in a treatment village in a low-saturation sublocation.

3. Pooled ineligible spillover regression: Effect on ineligibles of residing in a treatment village in high-saturation and low-saturation sublocations. This is our main analysis of interest for spillovers.

3.3 Interaction effects

To estimate interactions between saturation and spillover effects, we calculate the following effects from equation 1, clustering at the saturation group level:

1. Effect of saturation on direct treatment effect:

$$[E[y_{ihvs,t=1}|T=1|E=1,H=1] - E[y_{ihvs,t=1}|T=0,E=1,H=1]] -$$
$$[E[y_{ihvs,t=1}|T=1|E=1,H=0] - E[y_{ihvs,t=1}|T=0,E=1,H=0]] =$$
$$[\beta_1 + \beta_4 + \beta_5 + \beta_7] - [\beta_1 + \beta_4] = \beta_5 + \beta_7$$

2. Effect of saturation on spillover effect:
$$[E[y_{ihvs,t=1}|T=1|E=0,H=1] - E[y_{ihvs,t=1}|T=0,E=0,H=1]] -$$
$$[E[y_{ihvs,t=1}|T=1|E=0,H=0] - E[y_{ihvs,t=1}|T=0,E=0,H=0]] =$$
$$[\beta_1 + \beta_5] - [\beta_1] = \beta_5$$

3.4   Estimating the spatial horizon of spillover effects

The analyses outlined so far assume that spillover effects occur only within sublocations. However, it is possible that spillovers may occur not at the level of administrative units, but at the spatial level, i.e. based on the distance of a given village from other treated or untreated villages, regardless of whether or not these villages are in the same or different sublocations. We therefore adopt a similar approach of estimating radius-based measures for household and enterprise outcomes as in the midline, household and GE PAPs. We consider the total amount distributed per capita, $Amt_i^{RR'}$, within a series of donuts with inner radius R and outer radius R' around household i. As random assignment does not guarantee that the Amt variables are exogenous as they also depend on the share of households in a given donut that are eligible, we therefore instrument for this quantity with the proportion of eligibile households in the same donut whose villages were assigned to treatment. This analysis will be conducted separately a) for eligible households in control villages; b) ineligible households in treatment and control villages; and c) all households. We examine distance ranges from 0-1km up to 9-10km. We then estimate the following model for both enterprise and household data, with $R' = R + 1$:

$$y_{ivs} = \sum_{R=0km}^{9km} \gamma_{RR'} \widetilde{Amt}_i^{RR'} + \varepsilon_{ivs}. \qquad (3)$$

Here, y_{ivs} is the outcome of interest for household i in village v in sublocation s, and $\widetilde{Amt}_i^{RR'}$ is the per-capita amount transferred to households within a given radius band relative to unit i, instrumented as described above. We use the Schwartz Bayesian Information Criterion to select the nested model with the optimal number of radius terms, while imposing weak monotonicity [see Haushofer et al., 2016, for full details]. Standard errors

are calculated via Conley [1999, 2008] using a uniform kernel up to the maximum radius of the nested model. We then test i) whether the $\gamma_{RR'}$ terms are jointly different from zero and ii) whether the $\gamma_{RR'}$ terms are equal to one another.

One potential concern with this strategy is that part of the variation in the treatment density around a location comes from the assignment of sublocations to high- and low-saturation. As a robustness check, we will estimate models that fully interact an indicator for being in a high saturation sublocation with the amount variables, and will present estimates with and without the restriction that the effects are the same.

3.5 Treatment effects over time

To analyze the development of our treatment effects over time, we use the randomized timing of the endline survey to assess whether treatment effects vary with the time elapsed since transfers to a given village or sublocation were made. To determine the treatment timing, we use the approach for assigning transfers to months outlined in the midline PAP [see Haushofer et al., 2016, for full details]. We assign all villages to a quartile based on the time elapsed between the delivery of transfers and the time of the endline at the village level. This results in four types of villages: those in which the interval between transfer delivery and endline was less than 16 months; those in which it was 16–19 months; those in which it was 19–22 months; and those in which it was greater than 22 months. We then repeat the above analyses for "timing quartiles", i.e. we run the same regressions separately for the first through fourth quartiles of villages and sublocations in the order in which they received treatment. Where required (e.g. for computation of the fiscal multiplier), the resulting estimates are summed to create an impulse response function.

3.6 Multiple inference correction

Multiple comparisons adjustment will be used as described in the household PAP. Specifically, we use the false discovery rate to adjust across the 10 main outcomes separately for each individual effect discussed above. In addition, we adjust across the different variables within each outcome family for each of these specifications when analyzing these variables in detail. Further details are given in the original household PAP. We note that this approach is only appropriate for the specific purpose of testing if any of the outcome variables reach statistical significance while ensuring that the overall rate of type I errors is α (conventionally, 0.01, 0.05, and 0.1). Readers interested in specific outcomes should

focus on the naïve (per-comparison) p-values associated with that outcome. In addition, different theories may suggest specific groupings of outcomes; for instance, those interested in health outcomes may wish to control for the false discovery rate within that group of outcomes. Thus, while we present the same FDR adjustment as in the household PAP for consistency, we stress that readers should focus on the FDR or naïve p-values depending on their objective.

References

Timothy G. Conley. GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45, September 1999. URL `https://ideas.repec.org/a/eee/econom/v92y1999i1p1-45.html`. 3.4

Timothy G. Conley. Spatial econometrics. In Steven N. Durlauf and Lawrence E. Blume, editors, *The New Palgrave Dictionary of Economics*, volume 7, pages 741–747. Palgrave Macmillan, Houndsmills, second edition edition, 2008. 3.4

Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker. Pre-analysis plan for midline data: General equilibrium effects of cash transfers. May 2016. 3.4, 3.5

Johannes Haushofer, Edward Miguel, Paul Niehaus, and Michael Walker. Ge effects of cash transfers: Pre-analysis plan for household welfare analysis. July 2017. URL `https://www.socialscienceregistry.org/trials/505`. 1, 2

David McKenzie. Beyond baseline and follow-up: The case for more T in experiments. *Journal of Development Economics*, 99(2):210–221, 2012. URL `http://ideas.repec.org/a/eee/deveco/v99y2012i2p210-221.html`. 3.1