Evaluating Indonesia's Unconditional Cash Transfer Program, 2005-6*

Samuel Bazzi† Sudarno Sumarto‡ Asep Suryahadi§

October 2012

3IE—Evaluation Report

Abstract

Targeted cash transfer programs have been an important policy tool in developing countries. This paper considers (i) how the timing of transfers affect household expenditure and labor supply responses, and (ii) how household expectations shape our interpretation of those responses. We study these issues in the context of a short-term program that provided quarterly unconditional transfers of 30 USD to over 19 million households in Indonesia. Our empirical strategy relies on nationally representative panel data, difference-in-difference reweighting estimators, and the staggered rollout of the second quarterly transfer. On average, beneficiary households that received the full two transfers by early 2006 do not differ from comparable non-beneficiaries in terms of per capita expenditure growth and changes in labor supply per adult. However, beneficiaries still awaiting their second transfer report 7 percentage point lower expenditure growth and a reduction in labor supply by an additional 1.5 hours per adult per week on average. The expenditure differences dissipate by early 2007, several months after the final transfer was received by all beneficiaries. We also exploit variation in transfers per capita to identify a small marginal propensity to consume out of transfer income (around 0.10). We reconcile the empirical results with the predictions of a simple permanent income model, consider rival (missing) data-driven explanations, and document similar household responses to other transitory changes in income.

*We gratefully acknowledge financial support from the International Initiative for Impact Evaluation (3ie) and thank the Central Bureau of Statistics (BPS) of Indonesia for providing data. Umbu Raya provided excellent research assistance. We thank Michael Clemens, Craig McIntosh, Paul Niehaus, Julia Tobias, and seminar participants at UC-San Diego and the University of Western Australia for useful discussion. We also thank Robert Sparrow for assistance in matching households across survey waves. Any errors that remain are exclusively ours.

†Department of Economics, University of California, San Diego. Email: sbazzi@ucsd.edu.

‡The SMERU Research Institute, Jakarta, Indonesia. Email: ssumarto@smeru.or.id.

§The SMERU Research Institute, Jakarta, Indonesia. Email: suryahadi@smeru.or.id.

1 Introduction

Targeted cash transfer (CT) programs have been an important policy tool in developing countries over the last decade. Transfer programs in numerous settings have been shown to improve education, health, and other welfare outcomes among poor households.[1] Recent studies have explored optimal program design along several dimensions (e.g., Baird et al., 2011; Barrera-Osorio et al., 2011; Carrillo and Ponce Jarrín, 2009; De Janvry and Sadoulet, 2006; Filmer and Schady, 2011). However, despite a rich and growing evaluation literature, we have limited evidence on (i) how the timing of these changes in nonlabor income may affect household expenditure and labor supply responses, and moreover (ii) how household expectations shape our interpretation of those responses.

In this paper, we investigate these issues in the context of a large-scale, short-term unconditional cash transfer (UCT) program in Indonesia. After slashing regressive fuel subsidies, the government of Indonesia provided nearly 19 million households with quarterly transfers of around 30 USD—roughly 1/8th of average quarterly household expenditures at baseline—between October 2005 and September 2006. We use well-timed, nationally representative household-level panel data from the National Socioeconomic Survey, *Susenas* to identify the household expenditure and labor supply response over two time horizons: (i) a short-term period after which beneficiary households had received one or two quarterly transfers and (ii) a medium-term period by which time the program had ceased.

Our identification strategy relies on multiple sources of variation in transfer income. First, we use a difference-in-difference (DID) procedure reweighting all households by their predicted probability of treatment (see Heckman et al., 1998; Abadie, 2005). Although beneficiaries were identified through a quasi-means testing process, we show that it is not possible to reconstruct the proxy means scores in a reliable enough manner to justify a fuzzy regression-discontinuity design. Nevertheless, our augmented model for predicting program receipt captures substantial variation in treatment status across households, and reweighting effectively rebalances UCT recipient and non-recipient households along baseline characteristics.

Second, we exploit the staggered rollout of the second transfer payment. The staggering arose as a result of delays in the local disbursement schedule across subdistricts.[2] Because the timing of the midline follow-up survey varied with respect to this schedule, we are able to identify variation in the timing (and hence cumulative magnitude) of transfers received across beneficiary households: one-fourth of all recipients were still awaiting their second transfers at the time of enumeration in early 2006. As we show, the staggering occurred primarily across large regions rather than across households within regions. Moreover, we document that the staggering process was as good as random insomuch as the timing of transfers and surveys cannot be explained by observable differences across regions in terms of remoteness, weather shocks, or level of development.

Lastly, because the size of the transfer (per disbursement) was fixed regardless of household size, the scale of benefits varied considerably across recipient households. We take several steps to show that the variation in transfers *per capita* is plausibly exogenous and hence can be used to identify an intensive margin treatment

[1] See Hanlon et al. (2010) for a comprehensive review of the literature.

[2] Indonesia's administrative divisions proceed from province to district to subdistrict to village. Local post offices at the subdistrict level were responsible for disbursing quarterly cash transfers to all villages within their jurisdiction.

effect.[3] De Janvry and Sadoulet (2006) make use of similar variation in treatment intensity imposed by the cap on total transfers in the *Progresa* program in Mexico, and Kaboski and Townsend (2005, 2011b) analogously exploit variation in fixed financial transfers across Thai villages that vary in population size.

Our first key empirical results suggest that the timing of transfer disbursements sharply affect the expenditure and labor supply response. We find no mean differences in household per capita expenditure growth between control households (i.e., comparable non-beneficiaries) and UCT beneficiaries that had received the full two transfers as expected by the time of follow-up enumeration in early 2006. However, the one-quarter of UCT recipients still awaiting their second transfer at the time of follow-up enumeration report seven percentage point lower expenditure growth on average than both control households and UCT recipients that had already received the second transfer. The largest differences across groups are found for food rather than non-food expenditures. This relatively large differential treatment effect dissipates by early 2007, several months after the final quarterly transfer was received by all beneficiaries. Despite the null treatment effects on expenditures, we do find that UCT benefits are associated with some differential movement out of—albeit also into—poverty over both the short- and medium-term horizons.

We also find similar differential treatment effects in terms of labor supply. Timely receipt of the second transfer by early 2006 has null effects on labor supply per adult,[4] but delayed receipt is associated with a decline of 1.5 hours per adult per week on average. Unlike the expenditure response, however, some of this difference across the two groups of beneficiaries persists through early 2007, suggesting potential adverse medium-run labor market effects.

These baseline estimates hold up to a battery of robustness checks as well as alternative double-robust and control function estimators (see Busso et al., forthcoming; Imbens and Wooldridge, 2009). However, pooling the two treatment groups (i.e., ignoring the delayed receipt of the second transfer) and estimating a conventional binary treatment effect would have understated both the labor and expenditure response to the program. Although the sharp differences in outcomes across treatment groups are relatively well-identified by the exogenous staggering, the comparisons with the pure control group require stronger assumptions. Identification hinges on the probability of receiving any transfers being orthogonal to time-varying unobservable determinants of expenditures (or labor supply). In other words, targeting agents must not have allocated eligibility on the basis of idiosyncratic shocks occurring between enumeration at baseline and follow-up but not captured in the latter survey.

Exploiting an identification strategy which does not hinge on these same assumptions, we find that the scale of the transfers also matters. Conditional on the observed differences across midline treatment levels, expenditure growth is increasing in transfers per capita. We estimate a marginal propensity to consume (MPC) out of UCT income of around 0.10. Although small, the estimated MPC is economically meaningful. An increase in household transfers per capita by 10 USD per quarter implies roughly a 5 percent increase in monthly household expenditures per capita. This within-treatment group comparison suggests an important intensive margin treatment effect, not unlike what has been found in other studies (e.g., Filmer and Schady, 2011).

[3]Baseline household sizes do not vary systematically across treatment (and control) groups. Nor does the UCT program have any effect on the change in household size between periods. Also, although 7 percent of recipients report obtaining less than the full 30 USD per disbursement, this added variation is uncorrelated with all observable household and village characteristics.

[4]Adults are defined as individuals 14 and older. We find no response of child labor to the program.

We attempt to rationalize the observed treatment effects of the UCT using a simple conceptual framework based on the permanent income hypothesis (PIH). In a canonical PIH model, household expenditures should exhibit a very small response to unanticipated transitory income shocks. The MPC out of UCT income should be equivalent to $r/(1+r)$, where r is the real interest rate. Insomuch as the UCT program was not yet conceived at baseline *Susenas* enumeration in early 2005, households observed at that time could not have anticipated their receipt (or not) of benefits beginning in late 2005. Thus, at midline in early 2006, households that received two transfers should have expenditure growth that is $[r/(1+r)] \times T$ larger on average than comparable non-recipients and $[r/(1+r)] \times T/2$ larger than households that received one transfer (where T is some proportional measure of transfers received).[5] Given that real interest rates $r \approx 0$ around this period,[6] it seems plausible that we could find a null midline average effect of the UCT program when comparing recipient expenditure growth with counterfactual non-recipients. Even with $r > 0$, any measurement error in expenditure or hours worked could make it difficult to obtain precise estimates of small treatment effects. These explanations could also generate the null effects observed at endline in early 2007.

However, the large reduction in expenditures among recipients that had only received one transfer by midline enumeration merits an alternative explanation—albeit one that is still grounded in the PIH. Suppose that (i) at the time of forecasting their short-term future income roughly six months prior to midline enumeration, all UCT beneficiaries fully anticipated their receipt of transfers by early 2006, and (ii) r is sufficiently greater than zero to generate mean observable differences across treatment levels. Condition (i) implies identical expenditure growth among control households and recipients that had obtained the full two transfers as expected by midline. This setup also implies relatively lower expenditure growth among UCT beneficiaries unexpectedly still awaiting their second quarterly transfer at follow-up. This gap can be interpreted as the effect of income falling short of forecasted expectations. In other words, savings were drawn were too early ahead of the (delayed) transfer. While neither expectations/timing convention is entirely dispositive, both highlight the value of the PIH in characterizing household responses to the UCT.[7]

We go on to show that Indonesian households exhibit similar expenditure responses to other transitory covariate income shocks. In particular, we find that household expenditures per capita in rural areas increase in response to unanticipated positive rainfall shocks. However, the excess sensitivity is only found among households engaged in agriculture, whereas the treatment effects did not differ along this dimension of heterogeneity.

Nevertheless, the null treatment effects of the UCT program seem to contradict results from numerous other settings. Why did unconditional cash transfers to relatively poor and (presumably) liquidity-constrained Indonesian households not yield the large expenditure gains typically found in the literature? One possibility raised in auxiliary fieldwork conducted by the authors (see Sumarto et al., 2006) is that UCT beneficiaries spent the transfer funds immediately within weeks if not days after receipt. If these expenditures took place sufficiently prior to enumeration, then the survey instrument might miss them. To the extent

[5]Note that these all-else-equal comparisons hinge on (i) the reweighting estimator balancing treatment and control households on baseline observable characteristics, and (ii) zero mean differences in unobservables across treatment and control households.

[6]Despite relatively high nominal interest rates, consumer prices of many goods were soaring around this time as a result of the fuel subsidy cutbacks and the initial impact of the ban on rice imports (see Bazzi, 2012).

[7]Of course, it is still possible that the differential treatment effects are entirely due to the difference in the amount of transfers received irrespective of timing. While we cannot rule this out entirely, the PIH could still explain the differential as being due to liquidity constraints that prevented transfer recipients from borrowing to smooth consumption between fully anticipated transfer dates.

that these funds were used on durables, this does not seem to be the case since the results are unchanged when using a pro-rated measure of durable expenditures over the past year (roughly, March 2005–March 2006) rather than the past month as in our baseline approach. Moreover, the results are robust to controlling for the date of midline enumeration in early 2006 (though we do not observe the date of transfer receipt). Yet, we cannot rule out that the funds were allocated entirely towards immediate food expenditures as these are only recorded in the week prior to survey enumeration. Ultimately, available survey data do not allow us to assess the actual amount of savings out of the transfer beyond that implied by our estimated MPC.[8]

Another possibility is that targeting agents effectively identified *ex ante* precisely those households that would likely experience the greatest adverse shock as a result of the generalized inflation caused by the fuel subsidy cutbacks. Although we account for a range of observable determinants of program participation, local program enumerators in mid-2005 surely relied upon a much larger information set than available to researchers in the baseline survey from early 2005. If such targeting based on expected negative welfare shocks took place systematically across Indonesia, then our estimates might simply reflect the differential forecasting ability of targeting agents.

Yet another possibility is that recipient households strategically underreported their expenditures so as to remain on beneficiary lists that were under public scrutiny at the time.[9] This source of non-classical measurement error could bias the treatment effects downward if recipients perceived their ongoing participation as being contingent on reported welfare levels. We attempt to test for this source of bias by controlling for whether the household was assigned to the initial list by the village head (potentially more prone to patronage) or by a regional government official outside the village (less prone to patronage). In doing so, we find little evidence of any differential treatment effects along this dimension of the program.

This paper offers new evidence on the importance of timing and expectations in understanding the effects of cash transfers on household behavior in low-income settings. There is a large literature examining household responses to transitory changes in nonlabor income in the United States (e.g. Hsieh, 2003; Sahm et al., 2012; Shapiro and Slemrod, 1995, 2009; Souleles, 1999). These studies draw similarly rich insights using the PIH to understand why, for example, consumers do not respond to anticipated changes in after-tax income. At the same time, the analogous connection with the PIH has not yet permeated the large literature on cash transfers in developing countries.

There are, however, a few important exceptions closely related to the present study. Bianchi and Bobba (forthcoming) show that conditional cash transfers (CCT) delivered through *Progresa* in Mexico increase entrepreneurial activity among beneficiaries in advance of their actually receiving the transfers. By exploiting the differential timing of the transfers across households, they are able to argue that the CCT increased entrepreneurship not only by relaxing liquidity constraints but also by encouraging risk-taking. Edmonds (2006) makes use of an analogous eligibility rule in the Child Support Grant program in South Africa, which gives rise to differences across beneficiary households in the timing of transfer receipt but not the eventual total amount received. Contrary to the predictions of a canonical PIH model, he finds that households reduce

[8] A related concern is that certain asset purchases go unreported. In the baseline results, we employ data from the short-form expenditure questionnaire, which includes broad categories of durable goods and asset purchases. However we find identical results in robustness checks using the long-form expenditure questionnaire, which contains much more detailed expenditure sub-categories and has been shown to yield higher total reported expenditures (see Pradhan, 2009).

[9] The first few months of the UCT program in 2005 generated a great deal of public controversy surrounding the allocation of benefits and widespread perception of mistargeting (see Cameron and Shah, 2012).

child labor and increase schooling in anticipation of future transfer income, attributing the result to binding liquidity constraints. Beyond these two reduced-form studies, a recent structural evaluation of *Progresa* make it possible to assess the effect of control households' expectations over future transfers on the observed treatment effects (Attanasio et al., 2012). Failing to account for such expectations can lead researchers to understate the magnitude of actual treatment effects. Although this bias did not arise in the case of *Progresa*, the insights raised by Attanasio et al. resonate with our findings, which suggest that failing to account for unmet expectations over the timing of transfers would have led to substantially understated (and even negative) treatment effects on expenditure and labor supply outcomes.

The remainder of the paper proceeds as follows. Section 2 provides background on the program and the dataset employed in the analysis; Section 3 motivates the evaluation model and details the identification strategy; Section 4 presents the primary empirical results; Section 5 reconciles the main empirical findings with insights from a simple permanent income model; and Section 6 concludes.

2 Cutbacks and Cash

In the midst of escalating global oil and gas prices in 2005, the Government of Indonesia (GoI) slashed fuel subsidies, raising regulated prices by a weighted average of 29 percent in February and then again by 114 percent in September. The measures yielded over 10 billion USD in annualized budgetary savings, a portion of which the GOI put towards the country's first large-scale unconditional cash transfer (UCT) program. This section details various features of the program relevant to understanding the household expenditure and labor supply response to the transfer income.

2.1 Fuel Subsidy Removal and Price Shocks

The subsidy reform proceeded in two stages. In March 2005, the government raised gasoline and automotive diesel prices by 33 and 27 percent respectively. After several months and some publicity, the GOI dramatically slashed subsidies on October 1st, effectively raising prices of the three fuel products by a weighted average of 114 percent. Previously immune to policy change, kerosene prices nearly tripled increasing by 186 percent while gasoline and diesel prices grew another 88 and 105 percent respectively.

The direct effect of these price shocks on household welfare would depend first and foremost on the incidence of fuel consumption. Based on nationally representative household survey (*Susenas*) data from February 2004 prior to the first round of subsidy downgrades, over 95 percent of Indonesian households consume at least one of the three main fuel products, and over 90 percent consume kerosene. Figure 1 examines the distribution of national fuel expenditures across deciles of household expenditure per capita in 2004. Automotive diesel and gasoline subsidies are most regressive while the overall incidence of kerosene consumption tends to be relatively flat across the distribution of income. The slight dip in kerosene consumption among wealthier households suggests a modest progressive element to kerosene subsidies.

Nevertheless, fuel products comprise a small share of overall household expenditures among both rich and poor. On average, the poorest decile of households allocate 3.7 percent of total monthly expenditures to kerosene while households in the richest decile spend only 1.9 percent. Nearly 93 percent of the poorest households and 80 percent of the richest households purchased kerosene in the month preceding enumeration in February 2004. Meanwhile, only 6 percent of the poorest households directly purchase gasoline compared to 46 percent of the richest households. The corresponding average budget shares for gasoline were 0.1 percent for the poorest households and 2.3 percent for the richest households. Therefore, although a large swathe of the population stood to be adversely affected by the kerosene and gasoline subsidy removals, these small budget shares suggest that the pass-through to purchasing power would have to occur through more indirect channels.

The regulated increase in fuel prices led to substantial consumer price inflation as a result of rising costs of transportation and production of goods with substantial fuel-based inputs. Over the period from February 2005 to February 2006, the CPI increased by 17.9%. Figure 2 shows the timing of the subsidy removals and the subsequent pass-through to other consumer goods and services. The year-on-year inflation rate provides a convenient benchmark as the nationally representative household surveys used in this study are conducted on an annual basis. The figure shows that the economy-wide effects from the limited downgrade of gasoline and diesel subsidies in early 2005 were relatively small compared to the large inflationary upswing brought on by the second round of cutbacks in late 2005. Also, the path of food prices appears to follow a similar trajectory as fuel prices, albeit for largely orthogonal reasons related to trade policy.[10]

2.2 UCT Program Implementation

With the fiscal savings generated by the subsidy cutbacks, the government implemented a targeted unconditional cash transfer (UCT) program beginning in October 2005 and culminating in September 2006. The stated goal of the program was to provide four quarterly disbursements of 300,000 Rupiah (Rp) (around 30 USD) to the poorest 30 percent of households beginning on October 1st. Political exigency would ultimately dictate targeting and implementation.

The targeting of beneficiaries proceeded in three stages. First, local government officials devised a large list of potential recipient households in August 2005 using a combination of own-discretion and community-based records from prior government programs . Second, using a minimalist survey instrument (known as PSE05.RT), the regional public statistical bureaus enumerated households on the initial list as well as others from additional government sources. The survey questions concerned: (1) floor type, (2) wall and roof type, (3) toilet facility, (4) electrical source, (5) cooking fuel source, (6) drinking water source, (7) frequency of meat consumption, (8) frequency of meal consumption, (9) frequency of purchase of new clothes, (10) access to public health facilities, (11) primary source of income, (12) educational attainment of household heads, (13) amount of savings and type of assets, and (14) floor width.[11] Lastly, the Central Statistics Bureau (BPS) used the survey data to implement a proxy-means test to generate the final list of eligible households by the end of September. Although the PSE05.RT data and PMT scores are not available, the baseline *Susenas* data, which we describe next, include close proxies for all questions except those concerning savings, assets, and frequency of consumption.

[10] Around late 2005, the price of domestically-produced rice—the main staple among the majority of Indonesian households—began a steep upward ascent due in small part to rising transport costs but mostly due to the government decision to ban rice imports in January 2004. While a boon to rice producers, the spike in rice prices had arguably more severe consequences for poor households than did the downsizing of fuel subsidies (see Simatupang and Timmer, 2008; McCulloch, 2008).

[11] In practice, only 35 percent of households report ever being visited by enumerators, and 8 percent did not know whether or not their household was visited (according to *Susenas* 2006). The majority of those enumerated were visited not by BPS officials but by local government officials.

3 Empirical Strategy

We employ several quasi-experimental identification strategies in order to evaluate the effect of the UCT on expenditure and labor supply outcomes. First, due to a rushed implementation schedule and weak pre-existing targeting infrastructure, many non-poor received benefits while many poor did not. These targeting errors prove useful for the purposes of constructing counterfactual non-recipient households through reweighting procedures. Second, for exogenous administrative reasons, the second quarterly transfer was staggered across regions with respect to the timing of the midline survey. We exploit this variation after showing that the staggering process is orthogonal to observable household and regional characteristics. Third, because all households received the same transfer amount per disbursement, we observe considerable variation in transfers *per capita*. This allows us to identify—under certain testable assumptions—an intensive margin treatment effect as well as the (quasi-)marginal propensity to consume (MPC) out of transfer income. In the remainder of this section, we first describe the *Susenas* panel data and then detail the empirical strategies for exploiting the multiple sources of treatment variation.

3.1 Data

We use three waves of nationally representative panel data from the National Socioeconomic Survey (known as *Susenas*) collected in February-April 2005, 2006 and 2007. After matching households across the 2005 and 2006 rounds, we obtain a balanced panel of 9,048 households. We also observe a subset of households ($N = 7,016$) again in February-April 2007.[12] *Susenas* 2005 provides a good baseline as it was implemented prior to the announcement of the UCT program and the large-scale subsidy cutbacks in October (see Figure 2).

Taking advantage of the spatial mismatch in the timing of the midline survey and the rollout of UCT disbursements, we observe three levels of treatment denoted by the number of disbursements $D \in \{0, 1, 2\}$ received by the time of *Susenas* enumeration in February-April 2006. We observe 2,444 households in the treatment group ($D > 0$), but 639 of these households had only received a single disbursement at the time of enumeration while the remaining 1,805 households had received two disbursements.[13]

The UCT program was intended to reach all poor and near-poor households (below 1.2 times the official region-specific poverty line). Recipient households were indeed poorer on average than non-recipients in early 2005 prior to the UCT program (see Table 1). Yet, there was still evidence of potential (i) *leakage* of benefits as 37 percent of UCT recipients were in the top three national per-capita expenditure quintiles, and (ii) *undercoverage* as half of the lowest quintile did not receive any benefits. Figure 3 bears out these targeting results. Whether benefits were actually mis-targeted based on proxy means scores is unanswerable with existing data for reasons discussed below. Regardless, though, only 50 (39) percent of poor (near-poor) households received any transfers. In theory, the distributional overlap across groups, which is correlated with observable covariates $\mathbf{X}_h$, should make it easier to identify credible counterfactual non-recipients.

Despite the convenient panel setup, the *Susenas* data have an important limitation in that households only report a subset of the eligibility questions from the PSE.05 survey. We directly observe eight of the

[12] The baseline survey contains 10,574 households, while the follow-up in February 2006 contains 9,892 households. The February 2007 survey meanwhile contains more than 55,000 households, a subset of which were interviewed in the two preceding years. See the notes to Tables 1 for details on panel construction.

[13] Unfortunately, we do not observe the date on which households in the panel data received each of the disbursements.

fourteen eligibility indicators in the February 2005 baseline survey.[14] Among the most important questions unavailable in the *Susenas* survey are those concerning frequency of meal consumption, assets, and savings proxies. While it is not possible to obtain the actual household proxy mean scores that would allow us to implement a regression discontinuity design, we can use the available questions in *Susenas* coupled with the district-specific coefficients for each qualifying criteria to construct a quasi-PMT score.[15] However, as we show in Appendix A, these reconstructed scores (i) fail to produce any (even remotely fuzzy) discontinuities around the stipulated thresholds, and (ii) underperform our estimated propensity scores (see below) based on a richer set of household characteristics plausibly available to local enumerators and village officials.

3.2 Identification

In general, we are interested in the average treatment effect on the treated (ATT) of receiving d relative to s disbursements. Denoting this estimator by $\tau_{ds} \equiv \mathbb{E}[Y(d) - Y(s)|D = d]$, we aim to identify three parameters of interest $\boldsymbol{\tau} \equiv (\tau_{10},\ \tau_{20},\ \tau_{21})$ using the following difference-in-difference specification for the change in log consumption (or some measure of hours worked),

$$\Delta \ln C_{ht} = \kappa + \tau_{10}\mathbf{1}\{D_h > 0\} + \tau_{21}\mathbf{1}\{D_h = 2\} + \Delta\varepsilon_{ht}, \tag{1}$$

where $\tau_{20} \equiv \tau_{21} + \tau_{10}$. By taking differences, we remove all variation in the time-invariant determinants of expenditures across households. The conventional binary treatment effects estimator lumping together single and multiple disbursement recipients, $\mathbb{E}[Y(1) - Y(0)|D = 1]$, is a weighted average of τ_{10} and τ_{20} with the weights equivalent to the share of UCT recipients at each level of treatment $d \in \{1, 2\}$. Given data and policy constraints, our goal is to ensure that the comparison of outcomes across groups is as close as possible to what one would observe if treatment status D had been assigned randomly.

We pursue a reweighting approach in which the contribution of non-recipient households to the counterfactual is directly proportional to their estimated odds of treatment, $\omega = \widehat{P}/(1 - \widehat{P})$, where $\widehat{P}$ is the household's predicted probability of receiving *any* UCT benefits. We estimate this propensity score as a saturated function of (i) all underlying components of the proxy means scores (available to us in *Susenas*), and (ii) additional household characteristics that would have been known to local (informal) targeting agents at the time of eligibility designation. Figure 4 demonstrates the substantial overlap in propensity scores for treatment ($D > 0$) and control ($D = 0$) households. The full set of underlying parameter estimates are reported in Table 2.[16] Given the considerable overlap, we then use the ω terms as inverse probability weights in order to rebalance recipient ($D > 0$) and non-recipient ($D = 0$) households along observable dimensions. Empirically, less than 5 percent of the covariates in Table 2 exhibit statistically significant mean differences (in t-tests) across recipients and non-recipients after re-weighting by ω. Under the assumption that there are no time-varying unobservable determinants of consumption growth correlated with UCT receipt, we can then interpret the conventional binary treatment effect causally (see Abadie, 2005).

[14] A second limitation is that the data structure pose a somewhat nonstandard attrition problem. Although attritors appear much more similar to non-recipients than recipients (see Table 1), we do not know which attritors between 2005 and 2006 actually received the UCT. We observe recipient status among the 2,034 attritors between 2006 and 2007, and somewhat reassuringly the ratio of recipients to non-recipients remains essentially unchanged across years. The attrition is largely attributable to the panel survey design, which drops and replaces around 20 percent of the original households at each new wave. Although inter-survey attrition is potentially a non-negligible problem, we ignore its consequences in the econometric results presented below. Nevertheless, all results are robust to reweighting the sample so as to account for the probability of attrition as a function of all observable characteristics used to predict treatment.

[15] We are grateful to Lisa Cameron and Hamonangan Ritonga for providing the PMT coefficients.

[16] Further details on the underlying variables and estimating equation can be found in Appendix A.

However, in order to identify the multivalued treatment effects in equation (1), we must (minimally) verify the exogeneity of the staggered rollout of the second quarterly disbursement. In Table 3, we show that the probability of receiving disbursement two conditional on receiving disbursement one, $\mathbb{P}(D = 2 \mid D > 0)$, is explained almost entirely by geographic fixed effects. Whereas household-level characteristics explain considerable variation in the probability of receiving any disbursements, $\mathbb{P}(D > 0)$, even with > 600 subdistrict fixed effects, household-level characteristics explain little variation in $\mathbb{P}(D = 2 \mid D > 0)$ after controlling for district or subdistrict fixed effects. The R-squared and F tests in columns 7-12 suggest that the staggering occurs largely across (sub)districts and is plausibly exogenous with respect to baseline household characteristics. This is reassuring given that the subdistricts each have a respective post office branch, which was responsible for disbursing the quarterly cash transfers.

Moreover, in Table 4, we show that geographic characteristics—both fixed (e.g., distance to urban centers) and time-varying (e.g., rainfall shocks)—cannot explain the spatial variation in staggering. There is little evidence that relatively poorer remote regions received the second disbursement any later than relatively wealthier, more central regions. Lastly, in results available upon request, we find that the date of survey enumeration in early 2006 is orthogonal to the level of treatment. In other words, households waiting for their second disbursement at the time of enumeration were not simply residing in regions enumerated at later dates. Thus, we are confident that the staggering process occurred for largely exogenous administrative reasons and hence can be used to identify multiple levels of treatment in the midline survey enumerated in early 2006.[17]

In Figure 5, we compare the distribution of log baseline household expenditures per capita across treatment levels. Given the exogeneity of the staggering process, it is not surprising to find that the distributions for treatment groups $D = 1$ and $D = 2$ are nearly identical and, in fact, statistically indistinguishable. Although mis-targeting was rife, the control group is still substantially richer at baseline than the treatment groups. However, once we reweight control households using ω, the control group distribution shifts leftward and overlaps with the treatment group distributions quite strongly. The slight disproportion of control households in the right tail of the distribution leads to a small albeit statistically significant mean difference across the treatment and control groups (at the 10% level). This slight imbalance in the baseline outcome in levels poses a potential source of bias but only insomuch as that imbalance cannot be explained by observable time-invariant determinants of consumption. Otherwise the first differences will remove any bias. Other baseline covariates are effectively balanced after reweighting by ω (results available upon request).

In addition to variation in the timing of the second quarterly, we also utilize the fixed transfer size to identify the marginal effect of an increase in transfers *per capita*. Baseline household sizes do not vary systematically across treatment (and control) groups. Figure 6 plots the distribution of transfers per capita (at midline in early 2006) for all recipients demonstrating the variation across households conditional on the number of disbursements d. The two disbursement recipients obtained median transfers per capita of Rp 150,000 (mean 179,000 Rp), and single disbursement recipients Rp 75,000 (mean 91,000 Rp). To identify the

[17]We do not consider other approaches to identifying multivalued treatment effects (see Imbens, 2000; Cattaneo, 2010) since the multivalued treatment in our case is plausibly exogenous with respect to household and geographic characteristics. Because the covariates determining binary treatment status have very little predictive power in distinguishing between individuals with one or two disbursements (see Table 3), the approaches for identifying multivalued treatment effects using the generalized propensity score (i.e., predicting multiple treatment levels) offer little advantage and introduce additional noise.

intensive margin treatment effects, we can simply augment equation (1) with observed $transfers/capita$ *and* an exhaustive set of indicators for household size.[18] Of course, this source of identification is not without caveats of its own. We address these in turn.

4 Empirical Results

Having shown (i) the baseline balance of reweighting control $D = 0$ households by their estimated odds of treatment and (ii) the plausible exogeneity of the staggering process, we present the main empirical results in this section. In what follows, we report estimates of the multivalued treatment effects $\boldsymbol{\tau}$. In addition to pure OLS, we consider four alternative reweighting estimators. All are predicated on the inverse probability weighting (IPW) approach. The double robust estimator augments the IPW specification with controls for the linear propensity scores ($\widehat{P}_h$) or the covariates ($\mathbf{X}_h$) used to predict those scores. The heterogeneous control function estimator introduces a fifth-order polynomial in the propensity scores and allows it to vary across recipients and non-recipients. A review of these estimators can be found in Busso et al. (2009) and Imbens and Wooldridge (2009). Following suggestions therein, we trim 38 households with $\hat{P}_h > \widetilde{p}$ where $\widetilde{p}$ is the optimal bound derived using the procedure in Crump et al. (2009). In all specifications, we also control for province fixed effects, which among other purposes, captures differential regional trends in (real) expenditure growth. Standard errors are clustered at the village level in keeping with the cluster-based sampling procedures of *Susenas*.

4.1 Expenditures

We begin by considering estimates of equation (1) for the log difference in consumption between t and $t+1$. The top panel in Table 5 presents our baseline results for the short-term period from 2005-6. We find a consistent pattern of differential treatment effects across all reweighting specifications discussed above: Recipients still awaiting their second disbursement at the time of enumeration in early 2006 have significantly lower expenditure growth—by roughly 7.5 percentage points—relative to non-recipients *and* recipient households with both disbursements. Moreover, recipients of two disbursements have identical expenditure growth as non-recipients.[19] These results are largely insensitive to the estimator used with the exception that the OLS estimates of τ_{10} and τ_{21} are slightly lower. However, had we pooled the two recipient groups and estimated a conventional binary treatment effect—essentially a weighted sum of τ_{10} and τ_{21} with the weights equal to one and the share of recipients with two transfers, respectively—we would have understated the expenditure gains to receiving the full two transfers as expected by early 2006.

Retaining the same specifications and moving ahead to 2007, the bottom panel of Table 5 shows that the differential treatment effects dissipate over the two-year time horizon. This is intuitive since the UCT program had terminated by the time of enumeration in February-April 2007, and all UCT recipients had received the full set of four quarterly disbursements.[20]

[18] Kaboski and Townsend (2011a) make use of analogous variation in village-level transfers per capita to identify the effect of the Thai Million Baht program, which allocated identical financial grants across villages of varying population size.

[19] Interestingly, the OLS estimates are statistically indistinguishable from the reweighting estimates. An optimistic interpretation of this similarity would be that selection bias is limited after taking first-differences of the dependent variable and hence may be largely confined to the cross-section. The less favorable reading would be that the reweighting approach (i.e., our estimated propensity scores) has not purged the sample of selection bias. Unfortunately, it is not possible to distinguish among these opposing alternatives.

[20] These estimates are not an artifact of the attrition of households between 2006 and 2007 survey rounds (see Section 3.1). Key results remain largely unaffected when reweighting the sample to account for the probability of attrition, which is unconditionally identical across treatment levels and largely an artifact of administrative randomness rather than systematic household or regional characteris-

The main findings in Table 5 hold up to a number of robustness checks:[21]

Timing of the Midline Survey

One concern with exploiting the staggered rollout is that we are merely picking up differences in the time at which households received *Susenas* enumerators. The identification strategy hinges on there being differences in the disbursement schedule across households observed at roughly identical points in time. To ensure that differential enumeration dates are not driving our results, we control for 65 distinct days of enumeration across the country. Doing so leaves the results unchanged.

Alternative "per capita" Formulations

Some authors argue that when looking at household expenditure outcomes, one should account for the fact that children require less consumption (particularly of food) than adults to attain equivalent levels of welfare (see Deaton, 1997; Olken, 2006). We allow for this possibility by treating children as 0.5 or 0.75 adult equivalents where children are aged 0-9 or 0-14 years. Again, the results are unchanged.

Regional Differences in Inflation

By including province fixed effects, we remove trend differences across regions in terms of inflation and hence of the passthrough from fuel price increases to other consumer goods. We take two additional steps to ensure that local price differences are not driving our results. First, we deflate nominal expenditures using the nearest of the fifty regional CPI measures. Second, we control for increases in the price of the goods basket used to construct the district-specific poverty lines.

Durable Goods Expenditures Beyond the Last Month

In the baseline regressions, we measure durable goods expenditures in the last month. In so doing, our measure of expenditures may have missed important purchases using UCT funds prior to January 2006. In other words, the UCT may have led to an increase in expenditures several months prior to midline enumeration and perhaps immediately after UCT receipt in October-December 2005. Hence our comparison of durable goods purchases in the early months of 2005 and 2006 might understate the large positive effects of the UCT had we compared those purchases going back over the full year prior to enumeration. This does not seem to be the case. Pro-rating annual non-food expenditures to the monthly level (or identically, pro-rating food expenditures to the annual level) leaves our key parameter estimates unchanged.

Alternative Geographic Fixed Effects and Clustering

All of the results in Table 5 are robust to including district fixed effects as well as to clustering standard errors at any administrative division above the village.[22]

Participation in Other Social Programs

Several other previously operative social programs continued alongside the UCT. Receipt of such programs might confound our estimates of τ parameters if, for example, the UCT disbursement schedule was timed

tics.

[21] Detailed tables for all of these robustness checks will be made available in an online appendix.

[22] Including subdistrict or village effects removes nearly all of the exogenous variation in the staggering of the second quarterly transfer and pushes the estimates closer to a simple binary treatment effects specification, which as noted earlier understates the expenditure response.

so as to reach those households lacking other programs first. We control for participation in other programs (including a rice subsidy scheme, scholarships for poor students, and subsidized health insurance for the poor) and the results remain similar to the baseline.

Systematic Underreporting of Expenditures

One concern is that in the midst of public scrutiny over perceived program leakage and undercoverage, UCT recipients and particularly those still awaiting their second disbursement systematically underreported their expenditures. This would lead to non-classical measurement error and could explain the null or negative treatment effects. We (partially) test for this by controlling for whether the household was assigned to the initial list by the village head (potentially more prone to patronage) or by a regional government official outside the village (less prone to patronage). Again, we find no systematic departures from the baseline findings.

Alternative Estimators for the Binary Treatment Effect

We also consider a range of alternative estimators for the binary treatment effect of receiving any UCT benefits including nearest-neighbor matching (Abadie and Imbens, 2005), local linear matching (Heckman et al., 1998), inverse probability tilting (IPT) (Graham et al., 2012), and quantile reweighting (Firpo, 2007). In all cases, the main qualitative and quantitative findings remain unchanged: the UCT did not yield additional expenditure growth beyond that reported by counterfactual non-recipients—either at the mean of the outcome or at different quantiles in the case of the Firpo (2007) estimator.

Decomposing Expenditure Growth

In Table 6, we find that the observed treatment effects are driven by differences in expenditures on food rather than non-food items. Using the most flexible, control function estimator, we cannot reject the null hypothesis that all three groups $d \in \{0, 1, 2\}$ have identical non-food expenditure growth. Over the medium term period 2005-2007, we find similar patterns with the minor exception that two disbursement recipients have slightly larger food expenditure growth than non-recipients ($\tau_{21} \approx 0.04$).

In Table 7, we further disaggregate food and non-food expenditure items. In keeping with the specification for aggregate expenditure growth, we restrict the estimates for each commodity group to those households with non-zero expenditures in both periods.[23] We find the same general pattern as with aggregate expenditure growth in Table 5. For most expenditure subcategories, recipients still awaiting their second disbursement have statistically significantly lower expenditure growth than recipients of two disbursements and non-recipients.[24] Moreover, the second disbursement almost entirely eliminates the gap between recipient and non-recipient expenditure growth. However, we do observe slightly lower growth in recipient expenditures on prepared foods and substantially higher growth (≈ 11.2 percentage points) on durable appliances. Other statistically precise differences are observed for non-staple food expenditures and transport/communications.

[23] The presence of zeros in one or both periods gives rise a panel data sample selection problem. A fully specified demand system is beyond the scope of the present study, and lacking instruments for the extensive margin, we focus on the intensive margin of expenditure growth.

[24] Two commodity groups, housing/utilities and debt/taxes, depart form this general pattern, though the results are statistically imprecise.

Intensive Margin Treatment Effects

Having found robust differential treatment effects according to the timing (and total magnitude) of transfers received, we now consider an additional source of variation in the intensive margin of treatment. In particular, we estimate the following equation:

$$\begin{aligned}\Delta \ln C_{ht} \quad = \quad & \kappa + \tau_{10}\mathbf{1}\{D_h > 0\} + \tau_{21}\mathbf{1}\{D_h = 2\} \\ & + \psi \; transfers/capita_h + \sum_{j=1}^{13} \beta_j \mathbf{1}\{HH\ size_h = j\} + \Delta\varepsilon_{ht},\end{aligned} \tag{2}$$

where (i) we retain the IPW reweighting strategy, and (ii) $transfers$ is the total amount of UCT funds (in 100,000s of Rupiah) received by enumeration in early 2006, and (iii) $capita$ and $HH\ size$ are household size. After removing (i) the multivalued treatment effects through reweighting and the disbursement indicators, and (ii) the independent effects of household size through β_j terms, all that remains is information on the scale (or intensity) of UCT benefits. Under the assumption that $\mathbb{E}[\Delta\varepsilon_{ht} HH\ size_h] = 0$ (after reweighting), ψ then identifies the marginal effect of an additional unit of non-labor income per capita.

Before considering estimates of equation 2, we address two potential concerns with the identification strategy underlying equation (2). First, if the UCT program caused changes in household size, then any observed effect on expenditures may reflect this intermediate relationship.[25] We rule this out in Table 8, which applies the same reweighing estimators to the difference in household size as the dependent variable.

Second, local officials in some regions extracted a portion of the officially mandated 300,000 Rp disbursement per beneficiary. Approximately 6.5 (8.5) percent of recipients were subject to these informal taxes at the time of obtaining their first (second) UCT disbursement.[26] If the incidence of informal taxes varied systematically across recipients depending on household size or other characteristics, then the estimated elasticity of outcome Y with respect to transfers per capita might be biased. In Table 9, we show that the probability of recipient household h being taxed is orthogonal to observable household characteristics. Tables 8 and 9 point to the plausible exogeneity of household size with respect to other variation of interest.

In Table 10, we report estimates of ψ from equation (2) for total, food, and non-food expenditures per capita. Columns 1-3 impose $\beta_j = 0$ for all j, and column 4 allows $\beta_j \neq 0\ \forall j$ to allow for unconditional scale effects in the growth in household expenditures/capita (e.g., larger households can better cope with shocks). The point estimates of 0.04-0.065 for total expenditures per capita imply a marginal propensity to consume (MPC) out of transfer income of around 0.08-0.11, where the MPC is simply the elasticity of expenditures per capita with respect to transfers per capita. The estimated MPC is slightly higher for non-food expenditures and when allowing for unconditional scale effects. Although small, these elasticities are economically meaningful. The estimates imply that an increase in household transfers per capita by 10 USD per quarter implies roughly a 5 percent increase in monthly expenditures per capita. We return to these estimates in Section 5 when discussing the theoretical implications.

[25]Note that controlling for the difference in household size does not solve the problem (see Angrist and Pischke, 2009, on the "bad" control problem).

[26]These taxes went primarily to officials in the village. According to recipients subjected to these taxes, the proceeds were meant to cover local ID/certificate administration, security at disbursement centers, but most were intended for redistribution to non-recipients deemed deserving by local officials. The portion allocated to supposed local redistribution increased from 40 percent at the first disbursement to 62 percent at the second disbursement. Among those taxed, the median amount also increased from 20,000 Rp to 50,000 Rp. These increases were likely due in part to the rising discontent with the initial eligibility lists.

Expenditure-Based Poverty Transitions

Before turning to labor supply results, we report in Table 11 the effects of the UCT program on changes in the poverty status of households. We estimate a multinomial logit equation with four possible outcomes: chronic poverty (i.e., poor in both periods t and $t+1$), moving into poverty (i.e., non-poor in t, poor in $t+1$), moving out of poverty (i.e., poor in t, non-poor in $t+1$), and never poor (i.e., non-poor in t and $t+1$). The Indonesian poverty lines are district-specific and are calculated separately for urban and rural areas based on a local food consumption basket relevant to relatively low-income households. When estimating the multinomial logit equation, we retain the flexible, control function reweighting specification as in earlier results.

The average marginal effects in Table 11 suggest that the UCT program had heterogeneous effects on poverty over the short-term period from early 2005 to early 2006. On the one hand, transfer recipients are more likely to stay poor and also become poor. Yet, we also observe that transfer receipt, particularly the second disbursement, is associated with movement out of poverty.

We observe similar patterns over the medium-term time horizon (in the bottom panel of Table 11) albeit with a few important exceptions. First, UCT benefits are associated with a large increase (of 0.13) in the probability of moving out of poverty. This holds regardless of the timing of second quarterly transfer at midline enumeration in 2006. Second, the correlation between UCT receipt and the probability of remaining poor falls by half relative to to the short-term time horizon. In sum, although the UCT benefits did not lead to dramatic increases in household expenditures, the program did enable some households to move out of (officially-defined) poverty over both the short- and medium-term.

4.2 Labor Supply

In this subsection, we briefly discuss the potential effects of the UCT on the labor supply of household members > 10 years old and not currently enrolled in school. (In results available upon request, we find no evidence that the UCT program led to changes in the labor supply of children enrolled in school.) Our preferred metric of labor supply is total hours worked per household divided by the number of working age adults not currently enrolled in school. We advocate this measure instead of a simple average over household members for several reasons. First, we wish to remain relatively agnostic as to the complex determinants of the intra-household substitutability of labor. Second, we aim to capture implicitly the dependency ratios for a given household. For example, if a certain household relies on the labor supply of two individuals, we would prefer to assign a larger increase in labor supply for a given hour compared with a household relying on the labor of three individuals. Third, lacking strong priors on functional form or a readily available instrumental variable, we avoid distinguishing between the extensive and intensive margins of labor force participation.

In Table 12, we consider the difference in labor supply between periods as the dependent variable and deploy the same set of reweighting estimators as before. We find that the first UCT disbursement is associated with a reduction of around 1.7 hours worked per adult in the last week. These are economically meaningful effects given a baseline mean of 22 hours worked per adult. However, they are not robust to the most flexible control function specification in column 5. Nor do there appear to be statistically meaningful differences between non-recipients and recipients that received two quarterly transfers by midline ($\hat{\tau}_{20}$ is null). In other words, the negative labor supply response in columns 2-4 is largely confined to those recipient households still awaiting their second transfer at the time of enumeration in early 2006. A potential explanation for this finding is that in early 2006 households had re-optimized their labor supply to a lower level in anticipation of receiving transfers at a given date in the near future. Insomuch as those decisions had persistent effects (e.g., previously declined positions were already filled), it may have been difficult for households to increase their labor in response to the delayed receipt of the second quarterly transfer.

The bottom panel of 12 shows that short term labor supply effects remain several months after the final disbursement arrived in late 2006. It is somewhat puzzling that these labor supply differentials persist well after the time by which all recipients should have received the full set of four quarterly transfers. One possibility is that the short-run persistence argument has long-run consequences.

Lastly, in Table 13, we show that although hours worked per adult are declining in transfers per capita (conditional on disbursements received), these effects are relatively small and statistically imprecise. Nor are there meaningful differences in the estimated effects over the short- versus medium-term.

5 Discussion

In this section, we reconcile the main empirical results (for expenditures) with a conceptual framework based on the permanent income hypothesis (PIH).

5.1 Interpreting Treatment Effects through the PIH

Starting from a standard Euler equation for household h in period t,

$$u'(C_{h,t-1}) = (1+\delta)^{-1}\mathbb{E}_{t-1}\left[(1+r)u'(C_{ht})\right],$$

the PIH under certainty equivalence (quadratic preferences, intertemporal separability, perfect credit markets) and income uncertainty implies

$$\Delta C_{ht} = \frac{r}{1+r}\left[1-\frac{1}{(1+r)^{T-t+1}}\right]^{-1}\sum_{\tau=0}^{T-t}(1+r)^{-\tau}(\mathbb{E}_t-\mathbb{E}_{t-1})Y_{h,t+\tau}, \tag{3}$$

where $Y_{h,t+\tau}=\varepsilon_{h,t+\tau}$ is income at time $t+\tau$ (see Jappelli and Pistaferri, 2010). Adding a permanent component to income $Y_{h,t+\tau}=P_{h,t+\tau}+\varepsilon_{h,t+\tau}$ (where $P_{h,t+\tau}=P_{h,t+\tau-1}+\upsilon_{ht}$) and pushing out to infinity, we obtain

$$\Delta C_{ht} = \frac{r}{1+r}\varepsilon_{ht}+\upsilon_{ht},$$

where period t savings is given by

$$S_{ht} = -\sum_{j=1}^{\infty}\frac{\mathbb{E}_t\Delta Y_{h,t+j}}{(1+r)^j} = \frac{1}{1+r}\varepsilon_{ht}. \tag{4}$$

These equations provide a simple framework for understanding the observed effects of the UCT program on consumption or expenditures.

In keeping with the empirical context, we consider expenditure growth between periods t (or $t+1$) and $t-1$ and abstract away from permanent components of income. Restating the above expressions in logs (after imposing the relevant assumptions on the utility function), equation (3) implies

$$\Delta \ln C_{ht} = \left(\frac{r}{1+r}\right)(\ln Y_{ht} - \mathbb{E}_{t-1} \ln Y_{ht}). \tag{5}$$

Suppose income $\ln Y_{ht} = W_{ht} + D_{ht}$ where W_{ht} is the real wage and D_{ht} is a potentially nonzero nominal government transfer, which by definition (and public law) is transitory. For simplicity, let $W_{ht} = \varepsilon_{ht}$.

Using equation (5), we consider several possibilities for the expenditure patterns of UCT recipients and non-recipients. First, consider non-recipients. Suppose that non-recipient household h' had no prior expectation of being a transfer beneficiary (i.e., they were informed at time $t-1$ that they would not be receiving any benefits in the future). This implies that their expenditure growth can be written as

$$\Delta \ln C_{h't} = \left(\frac{r}{1+r}\right)\varepsilon_{h't}. \tag{6}$$

There are now multiple cases to consider for UCT recipients. First, suppose that all identified beneficiaries anticipated (at time $t-1$) that they would have received two transfer disbursements by time t in early 2006. Then, for recipients that realized two transfer disbursements D by enumeration in early 2006, we obtain

$$\Delta \ln C_{ht} = \left(\frac{r}{1+r}\right)\varepsilon_{ht}. \tag{7}$$

That is, on average, these households exhibit identical expenditure growth to non-recipients. Empirically, the reweighting procedure detailed above ensures that recipient and non-recipient households draw from same income distribution (i.e., $\mathbb{E}[\varepsilon_{ht}] = \mathbb{E}[\varepsilon_{h't}]$). However, for those recipients that realized only one transfer by enumeration in early 2006,

$$\Delta \ln C_{ht} = \left(\frac{r}{1+r}\right)(\varepsilon_{ht} - D_{ht}), \tag{8}$$

where the $-D_{ht}$ term captures the "surprise" effect of not having received the second disbursement by the time anticipated ex ante. In other words, these households would have drawn down savings in anticipation of the second disbursement (see equation (4)). However, its late arrival meant that the household was left with insufficient liquidity in the week(s) just prior to *Susenas* enumeration.

If we define $t-1$ as the period immediately after the announcement of the program benefits and implementation schedule, then equations (6)-(8) provide a justification for the treatment effects reported in Table 5. These equations are also consistent with the largest expenditure differences being observed for food rather than non-food items (see Table 6) since the former is reported over the week immediately prior to enumeration whereas the latter is reported over the month prior to enumeration. Moreover, this framework can also explain why the differential treatment effects in Table 5 dissipate by 2007. Taking a longer two-period difference in log expenditures between $t-1$ and $t+1$, the surprise effect in equation (8) no longer holds as all recipients received all four quarterly transfers as expected by the end of 2006.

On the other hand, if eventual recipient households did not anticipate the UCT program at time $t-1$, then equation (5) implies

$$\Delta \ln C_{ht} = \left(\frac{r}{1+r}\right)(\varepsilon_{ht} + 2D_{ht}) \tag{9}$$

for households realizing two disbursements by early 2006 and

$$\Delta \ln C_{ht} = \left(\frac{r}{1+r}\right)(\varepsilon_{ht} + D_{ht}) \tag{10}$$

for households realizing only one disbursement. For $r > 0$, this implies (i) that recipients should have higher expenditure growth than non-recipients and (ii) that recipients of two disbursements should have higher growth than recipients of one disbursement. Implication (ii) is borne out in Table 5, but implication (i) is not.

Although both formulations are informative, neither is dispositive. Taking the period around baseline enumeration in February-April 2005 as $t-1$ (and the midline follow-up enumeration period in February-April 2006 as t), no Indonesian household could have anticipated the subsidy cuts and cash transfer program implemented later that year since the government had not yet publicized their plans for such a program. Given this timeline, it is difficult to justify the initial formulation despite its obvious appeal. Yet, the second formulation in which the transfers were entirely unexpected requires ignoring the sharp break in expectations over future income that occurred around September 2005 midway between baseline and midline follow-up enumeration. One can see the potential problems with this by taking sub-annual time horizons in equation (3) and recalling that food expenditures are reported over the last week while non-food expenditures are reported over the last month (or last year).

Despite their stark differences, the two timing regimes coincide in the predicted expenditure growth differential between recipients of one relative to two disbursements over the short term. Under both expectations regimes, recipients of two disbursements should have expenditure growth that is roughly $[r/(1+r)] \times D$ greater than recipients still awaiting their second disbursement at the time of enumeration. Taking the estimates of τ_{21} from Table 5, we obtain $\widehat{r} \approx 0.075$. Moreover, as discussed above, the estimates of ψ in Table 10 imply a marginal propensity to consume out of transfer of around 0.08-0.10, implying similar estimates of r.

In practice, all of the above predictions hinge on the real interest rate r being non-zero. While nominal interest rates quoted by the government were indeed quite high around this time, so was inflation on account of the fuel subsidy cutbacks. Even if $r \approx 0$, households may respond to transitory income shocks if they are liquidity-constrained because, for example, credit markets are imperfect. We turn now to a test of this prediction among Indonesian households in our sample in order to rule out concerns that the household response to the UCT was somehow anomalous and the PIH-based explanation spurious.

5.2 Household Responses to Other Transitory Income Shocks

Given evidence that Indonesian households respond to transitory UCT benefits, we examine in this brief subsection whether households exhibit similar responses to other types of transitory income shocks. Following others in the development literature beginning with Paxson (1992), we exploit spatial and time series variation in rainfall, a transitory source of income fluctuations across the Indonesian archipelago. For individual h residing in village v, we measure the transitory rainfall shock in year t as the log rainfall level in that district over the province-specific growing season minus the log mean rainfall level for that district over the forty years/seasons prior to t.[27]

[27] Due to merging difficulties, we are forced to drop households residing in villages on Papua.

The key message from Table 14 is that transitory rainfall shocks are associated with higher growth in household expenditures. However, the expenditure response is largely confined to those households in the agricultural sector, and particularly those with any land-holdings. In column 1, we find no relationship between rainfall shocks and consumption. However, allowing the elasticity to vary across rural and urban areas in column 2, we find a positive elasticity of expenditure growth with respect to rainfall shocks that is around 0.13 albeit imprecisely estimated. Moreover, in columns 3-5, we find that households reporting agricultural activities as their primary income and owning any agricultural land exhibit a small albeit statistically precise and economically meaningful expenditure response to transitory changes in income associated with rainfall shocks. The estimate in column 3 suggests that in agricultural households, a 10 percent deviation of rainfall from its long-run mean yields roughly a 2.2 percent increase in consumption. The elasticity is of similar magnitude in column 4 when restricting to land-owning households. Taken together, these results suggest that Indonesian household expenditures are more responsive to transitory income shocks than would be predicted under the classical permanent income hypothesis in the absence of borrowing constraints.[28]

6 Conclusion

This paper has considered the importance of timing and expectations in interpreting the household expenditure response to unconditional cash transfers in Indonesia. Our empirical strategy relied on nationally representative panel data, difference-in-difference reweighting estimators, and the staggered rollout of the second quarterly transfer. Our findings highlight the benefit of having multiple sources of variation in transfer income. The staggered rollout allowed us to identify differential treatment effects depending on the timing of the second transfer. On average, beneficiary households that received the full two transfers as expected by early 2006 do not differ from comparable non-beneficiaries in terms of per capita expenditure growth and changes in labor supply per adult. However, beneficiaries still unexpectedly awaiting their second transfer report 7 percentage point lower expenditure growth and a differential reduction in labor supply by an additional 1.5 hours per adult per week on average. Using the third wave of panel data, we find that the expenditure differences dissipate by early 2007, several months after the final transfer was received by all beneficiaries. Using the fact that the transfer amount per disbursement was fixed across households, we are able to identify a small, short-run marginal propensity to consume out of transfer income of around 0.10. We reconcile our findings with insights of a simple permanent income model and largely rule out alternative explanations based on missing or imperfect data.

In addition to offering a new way of framing the household response to unconditional cash transfers, our paper also relates more generally to the literature on the role of cash transfers in policy reform in developing countries. Unlike numerous programs in Latin America and elsewhere, the UCT in Indonesia was not explicitly designed as a transformative poverty alleviation program. Rather, the government used the program as means of transitioning away from regressive fuel subsidies. Similar subsidy reforms have either recently been implemented or are being considered across a number of developing countries (Coady et al., 2010). These programs have a number of important welfare implications and warrant further study. Our results from Indonesia suggest that the household response to cash transfers in such contexts may hinge strongly on perceived program duration as well as the timing of the transfers with respect to subsidy cutbacks.

[28] Although rainfall shocks only affect the transitory income and hence expenditures of certain segments of the (rural) population, the UCT benefits and especially the intensive margin of treatment do not have heterogeneous effects along these same dimensions. These results (available upon request) increase our confidence in the interpretation of the UCT benefits as a transitory income shock in the context of the PIH framework considered above.

References

A. Abadie. Semiparametric difference-in-differences estimators. *Review of Economic Studies*, 72(1):1–19, 2005.

A. Abadie and G.W. Imbens. Large sample properties of matching estimators for average treatment effects. *Econometrica*, 74(1):235–267, 2005.

J.D. Angrist and J.S. Pischke. *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press, 2009.

O.P. Attanasio, C. Meghir, and A. Santiago. Education choices in mexico: Using a structural model and a randomized experiment to evaluate progresa. *The Review of Economic Studies*, 79(1):37–66, 2012.

S. Baird, C. McIntosh, and B. Özler. Cash or condition? evidence from a cash transfer experiment. *The Quarterly Journal of Economics*, 126(4):1709–1753, 2011.

F. Barrera-Osorio, M. Bertrand, L.L. Linden, and F. Perez-Calle. Improving the design of conditional transfer programs: Evidence from a randomized education experiment in colombia. *American Economic Journal: Applied Economics*, 3(2):167–195, 2011.

S. Bazzi. Wealth heterogeneity, income shocks, and international migration:. *Unpublished Manuscript*, 2012.

M. Bianchi and M. Bobba. Liquidity, risk, and occupational choices. *Review of Economic Studies*, forthcoming.

M. Busso, J. DiNardo, and J. McCrary. New evidence on the finite sample properties of propensity score matching and reweighting estimators. *IZA Discussion Papers*, 2009.

M. Busso, J. DiNardo, and J. McCrary. Finite sample properties of semiparametric estimators of average treatment effects. *Journal of Business and Economic Statistics*, forthcoming.

L. Cameron and M. Shah. Can mistargeting destroy social capital and stimulate crime? evidence from a cash transfer program in indonesia. *Unpublished Manuscript*, 2012.

P.E. Carrillo and J. Ponce Jarrín. Efficient delivery of subsidies to the poor: Improving the design of a cash transfer program in ecuador. *Journal of Development Economics*, 90(2):276–284, 2009.

M.D. Cattaneo. Efficient semiparametric estimation of multi-valued treatment effects under ignorability. *Journal of Econometrics*, 155(2):138–154, 2010.

D. Coady, J. Tyson, J.M. Piotrowski, R. Gillingham, R. Ossowski, and S. Tareq. *Petroleum Product Subsidies: Costly, Inequitable, and On the Rise*. International Monetary Fund, 2010.

R.K. Crump, V.J. Hotz, G.W. Imbens, and O.A. Mitnik. Dealing with limited overlap in estimation of average treatment effects. *Biometrika*, 96(1):187–199, 2009.

A. De Janvry and E. Sadoulet. Making conditional cash transfer programs more efficient: designing for maximum effect of the conditionality. *The World Bank Economic Review*, 20(1):1–29, 2006.

A. Deaton. *The analysis of household surveys: A microeconometric approach to development policy*. Johns Hopkins University Press, 1997.

E.V. Edmonds. Child labor and schooling responses to anticipated income in south africa. *Journal of Development Economics*, 81(2):386–414, 2006.

D. Filmer and N. Schady. Does more cash in conditional cash transfer programs always lead to larger impacts on school attendance? *Journal of Development Economics*, 96(1):150–157, 2011.

S. Firpo. Efficient semiparametric estimation of quantile treatment effects. *Econometrica*, 75(1):259–276, 2007.

B.S. Graham, C.C.D.X. Pinto, and D. Egel. Inverse probability tilting for moment condition models with missing data. *The Review of Economic Studies*, 79(3):1053–1079, 2012.

J. Hanlon, D. Hulme, and A. Barrientos. *Just give money to the poor: The development revolution from the global South*. Kumarian Press, 2010.

J.J. Heckman, H. Ichimura, and P. Todd. Matching as an econometric evaluation estimator. *Review of Economic studies*, 65(2):261–294, 1998.

C.T. Hsieh. Do consumers react to anticipated income changes? evidence from the alaska permanent fund. *American Economic Review*, 93(1):397–405, 2003.

G.W. Imbens. The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3):706–710, 2000.

G.W. Imbens and J.M. Wooldridge. Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86, 2009.

T. Jappelli and L. Pistaferri. The consumption response to income changes. *Annual Review of Economics*, 2: 479–506, 2010.

J.P. Kaboski and R.M. Townsend. Policies and impact: An analysis of village-level microfinance institutions. *Journal of the European Economic Association*, 3(1):1–50, 2005.

J.P. Kaboski and R.M. Townsend. A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica*, 79(5):1357–1406, 2011a.

J.P. Kaboski and R.M. Townsend. A structural evaluation of a large-scale quasi-experimental microfinance initiative. *Econometrica*, 79(5):1357–1406, 2011b.

R.W. Klein and R.H. Spady. An efficient semiparametric estimator for binary response models. *Econometrica*, 61:387–421, 1993.

N. McCulloch. Rice prices and poverty in indonesia. *Bulletin of Indonesian Economic Studies*, 44(1):45–64, 2008.

B.A. Olken. Corruption and the costs of redistribution: Micro evidence from indonesia. *Journal of Public Economics*, 90(4):853–870, 2006.

C.H. Paxson. Using weather variability to estimate the response of savings to transitory income in thailand. *The American Economic Review*, 82:15–33, 1992.

M. Pradhan. Welfare analysis with a proxy consumption measure: Evidence from a repeated experiment in indonesia. *Fiscal Studies*, 30(3-4):391–417, 2009.

C.R. Sahm, M.D. Shapiro, and J. Slemrod. Check in the mail or more in the paycheck: Does the effectiveness of fiscal stimulus depend on how it is delivered? *American Economic Journal: Economic Policy*, 4(3):216–250, 2012.

M.D. Shapiro and J. Slemrod. Consumer response to the timing of income: Evidence from a change in tax withholding. *American Economic Review*, 85(1):274–283, 1995.

M.D. Shapiro and J. Slemrod. Did the 2008 tax rebates stimulate spending? *American Economic Review*, 99(2): 374–379, 2009.

P. Simatupang and C.P. Timmer. Indonesian rice production: Policies and realities. *Bulletin of Indonesian Economic Studies*, 44(1):65–80, 2008.

N.S. Souleles. The response of household consumption to income tax refunds. *The American Economic Review*, 89(4):947–958, 1999.

S. Sumarto, N. Toyamah, S. Usman, B. Sulaksono, S. Budiyati, W.D. Widyanti, M. Rosfadhila, H. Sadaly, S. Erlita, R.J. Sodo, and S. Bazzi. A rapid appraisal of the implementation of the 2005 direct cash transfer program in indonesia: A case study in five kabupaten/kota. *Development Economics Working Papers*, 2006.

Figures

Figure 1: Benefit Incidence of Fuel Subsidies, 2004

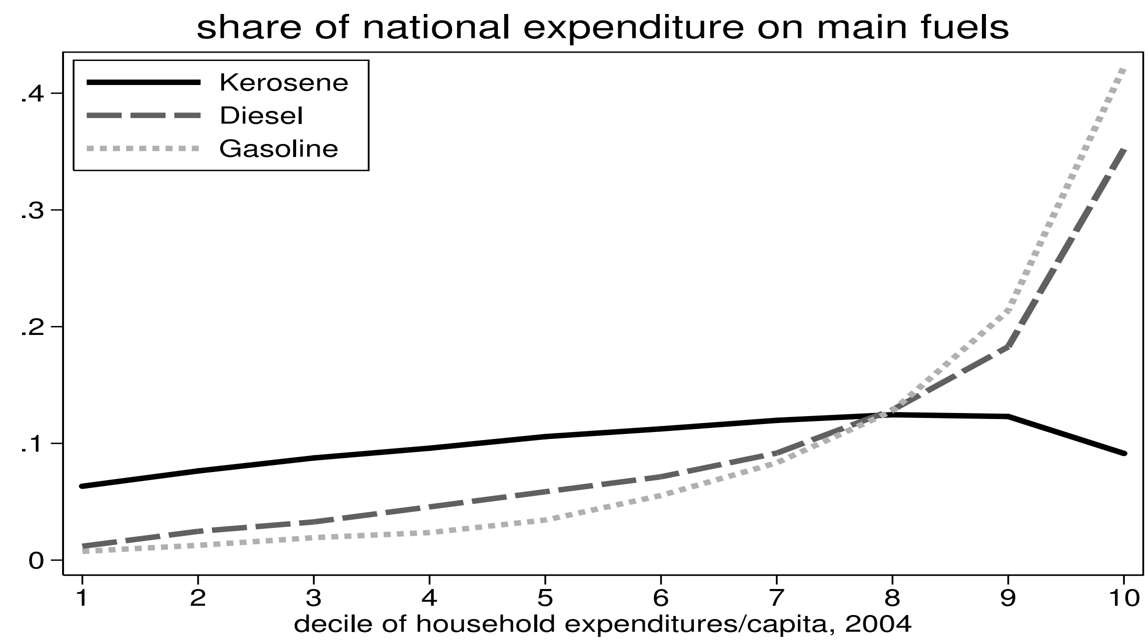

Notes: Calculated from *Susenas* 2004. Each point on the line represents that decile×location share of overall national expenditure on the given fuel product.

Figure 2: Subsidies, Transfers and Surveys: A Timeline of Events

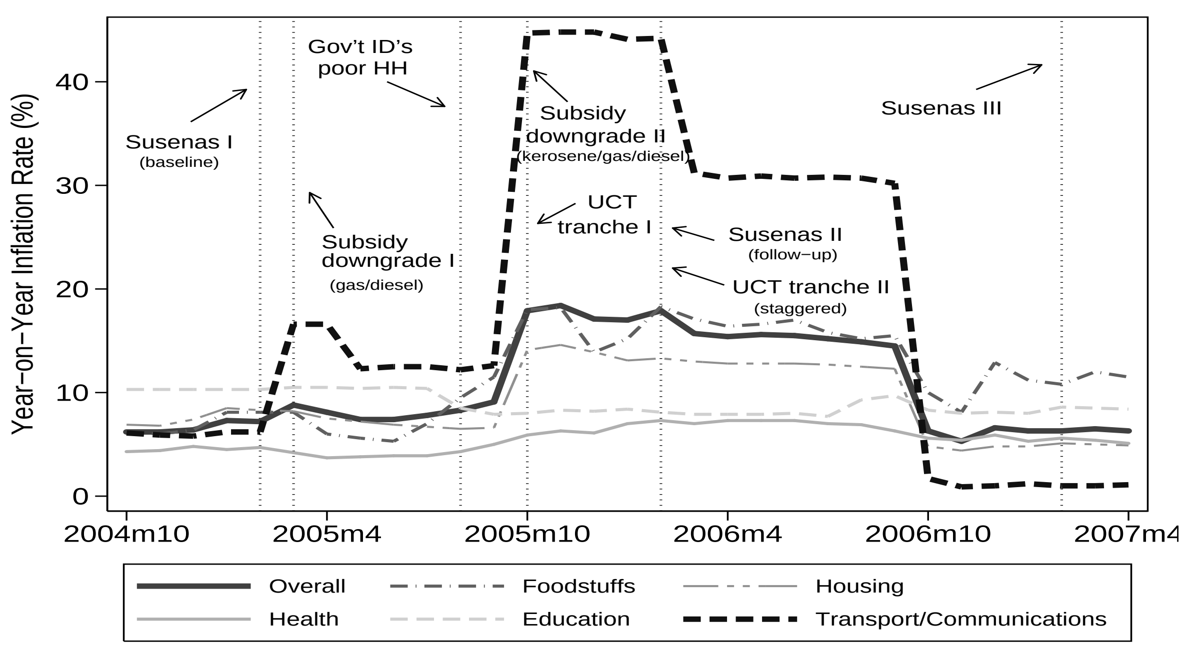

Notes: Monthly price indices are obtained from Bank Indonesia's online data system, November 2009.

Figure 3: Treatment Level by Baseline Expenditure Decile

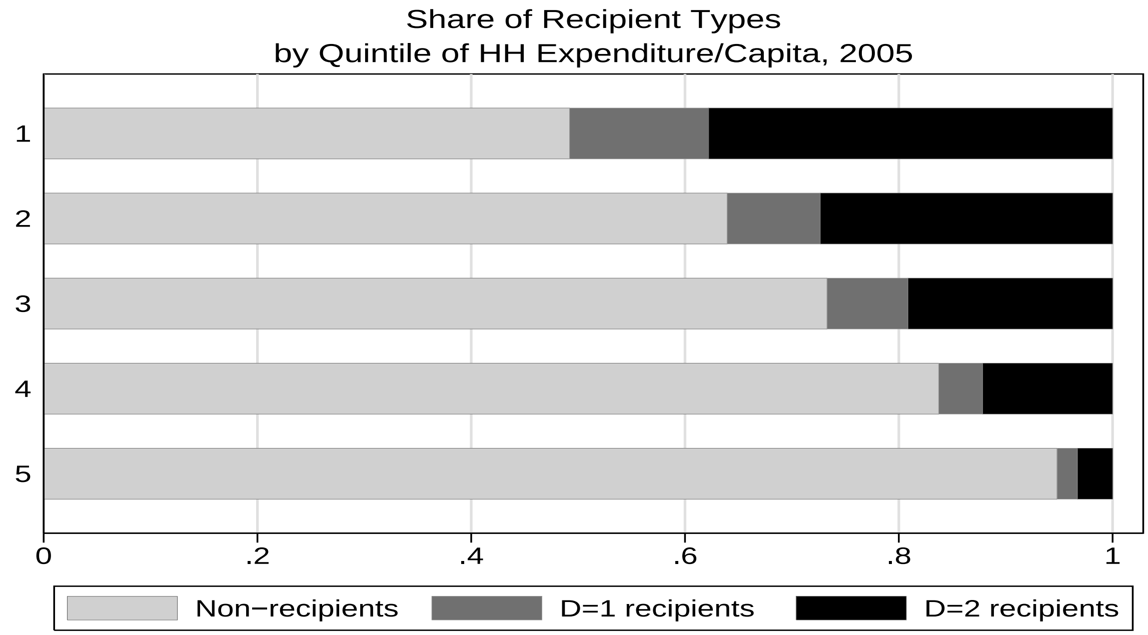

Notes: $D = d$ recipients obtained d UCT disbursements by enumeration in early 2006 as reported in a module attached to *Susenas* 2006. The quintile of household expenditures per capita is based on data reported in *Susenas* 2005.

Figure 4: Overlap in Estimated Propensity Scores ($\hat{P}$)

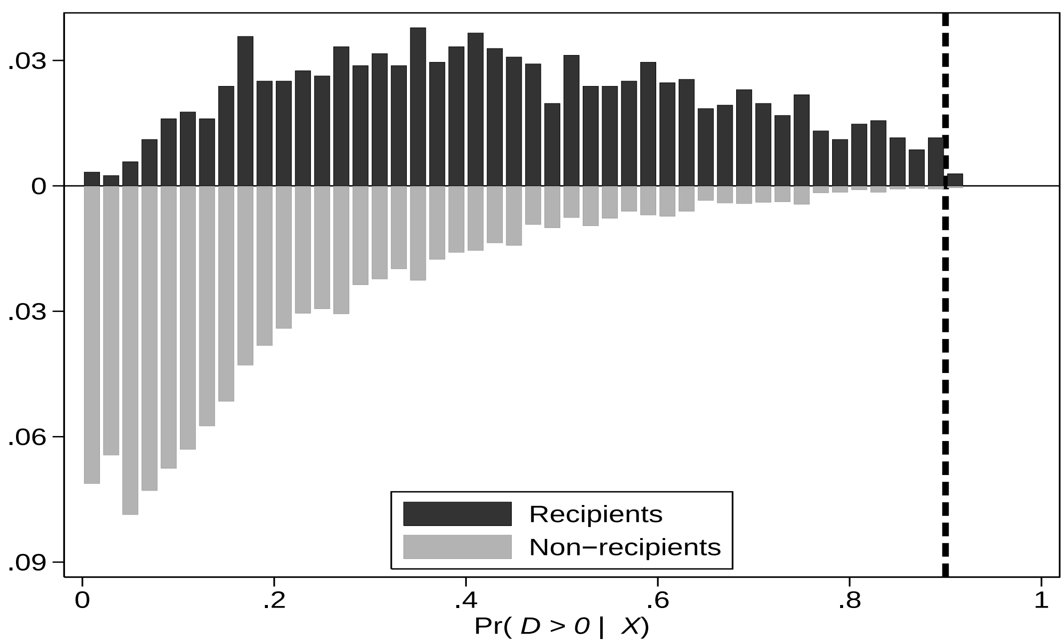

Notes: Propensity scores obtained from flexible logit regressions (see Table 2. Observations to the left of the dashed vertical line fall within the Crump et al. (2009) optimal overlap region.

Figure 5: Baseline Expenditure Distributions by Treatment Status

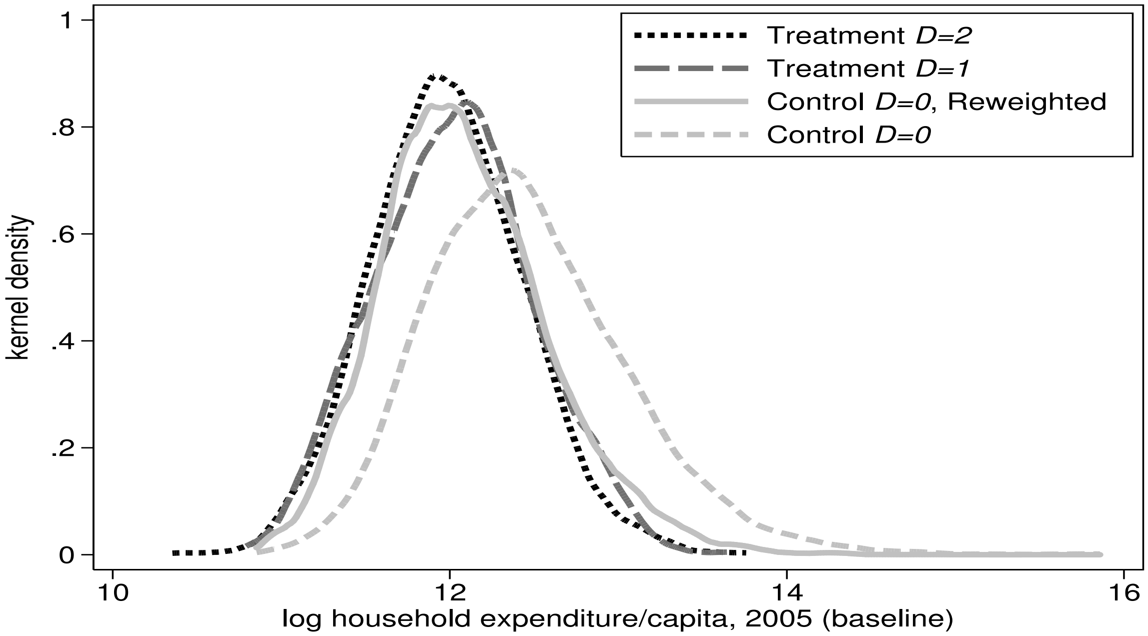

Notes: All distributions estimated using Epanechnikov kernel and a rule-of-thumb bandwidth. The "Control (Reweighted)" observations are adjusted using inverse probability weights (IPW) based on normalized estimated odds of treatment $\omega = \hat{P}/(1 - \hat{P})$.

Figure 6: Distribution of Transfers per Capita through February 2006

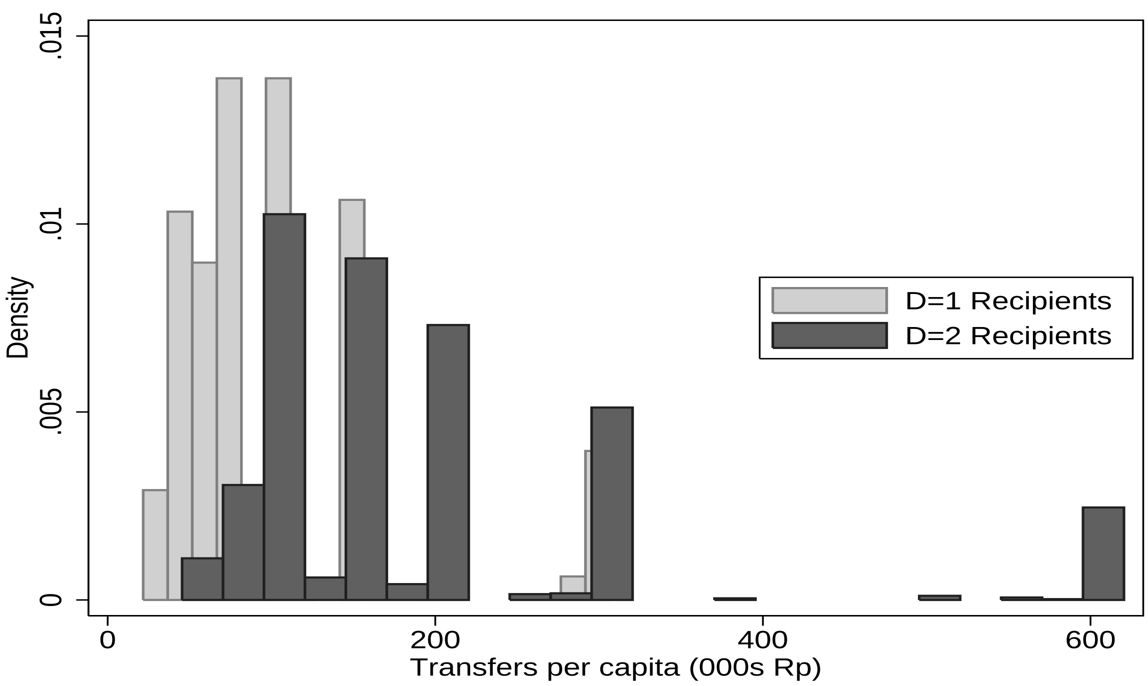

Notes: The transfer amount reported by households is obtained from a module attached to *Susenas* 2006.

Tables

Table 1: Expenditure statistics, 2005 and 2006

	2005					2006				
	Mean	SD	Min	Median	Max	Mean	SD	Min	Median	Max
	Non-recipients ($N = 6606$)									
Expenditure/capita (000s Rp)	315	292	52	243	7702	356	300	31	272	4891
Food expenditure/capita (000s Rp)	162	93	30	138	2790	182	104	20	155	1141
Non-food expenditure/capita (000s Rp)	153	234	8	94	7071	174	228	0	108	4236
Education expenditure/capita (000s Rp)	11	60	0	2	2269	8	41	0	0	1660
Health expenditure/capita (000s Rp)	11	67	0	2	2607	10	62	0	2	3137
Below poverty line	0.10	0.30	0	0	1	0.11	0.31	0	0	1
Quintile (nat'l) expenditure/capita	3.23	1.38	1	3	5	3.28	1.37	1	3	5
Quintile (intra-province) expenditure/capita	3.21	1.39	1	3	5	3.26	1.38	1	3	5
	$D = 1$ ***Recipients*** ($N = 639$)									
Expenditure/capita (000s Rp)	185	93	49	165	843	195	118	41	170	1817
Food expenditure/capita (000s Rp)	121	60	32	110	761	123	62	30	110	422
Non-food expenditure/capita (000s Rp)	65	49	9	52	423	72	80	9	56	1581
Education expenditure/capita (000s Rp)	2	6	0	0	220	2	4	0	0	48
Health expenditure/capita (000s Rp)	10	83	0	1	1832	4	11	0	1	150
Below poverty line	0.25	0.43	0	0	1	0.34	0.47	0	0	1
Quintile (nat'l) expenditure/capita	2.25	1.21	1	2	5	2.14	1.18	1	2	5
Quintile (intra-province) expenditure/capita	2.36	1.27	1	2	5	2.27	1.25	1	2	5
	$D = 2$ ***Recipients*** ($N = 1805$)									
Expenditure/capita (000s Rp)	178	90	31	159	945	192	92	37	172	908
Food expenditure/capita (000s Rp)	115	54	17	104	645	124	57	23	112	484
Non-food expenditure/capita (000s Rp)	63	50	9	50	576	68	51	0	55	682
Education expenditure/capita (000s Rp)	3	8	0	0.4	220	2	5	0	0	68
Health expenditure/capita (000s Rp)	5	11	0	2	178	5	22	0	1	751
Below poverty line	0.28	0.45	0	0	1	0.31	0.46	0	0	1
Quintile (nat'l) expenditure/capita	2.12	1.16	1	2	5	2.11	1.12	1	2	5
Quintile (intra-province) expenditure/capita	2.28	1.25	1	2	5	2.27	1.22	1	2	5
	Attritors ($N = 771$)									
Expenditure/capita (000s Rp)	323	272	54	252	2927					
Food expenditure/capita (000s Rp)	180	119	38	150	1073					
Non-food expenditure/capita (000s Rp)	142	197	10	86	2497					
Education expenditure/capita (000s Rp)	7	27	0	0.4	563					
Health expenditure/capita (000s Rp)	13	55	0	2	750					
Below poverty line	0.14	0.35	0	0	1					
Quintile (nat'l) expenditure/capita	3.23	1.41	1	3	5					
Quintile (intra-province) expenditure/capita	3.24	1.42	1	3	5					

Notes: A balanced two-year panel is constructed by matching along (a) province-district-subdistrict-village-sampling ID-household ID and (b) household head names in the 2005 and 2006 *Susenas* panels. While a traditional merge along strict geographic identifiers provides a balanced panel of $9,797$ households, significant discrepancies in household characteristics (including first names of household members) across waves indicate that survey administrators did not ensure the time-consistency of household presence in the physical location of prior enumeration. A name-matching algorithm provided by Robert Sparrow generated an initial balanced panel of almost 8500 households, and through further manual inspection, an additional 550 households were added. $D = d$ recipients obtained d UCT disbursements by enumeration in early 2006. Attritors are those households which could be identified in the 2005 baseline survey but not in the subsequent rounds. *Variable description: Rp* stands for *Rupiah*. The exchange fluctuated between $9,500$ Rp and $10,500$ Rp to the dollar between October 2005 and September 2006. All expenditure variables are household per capita expressed in *Rupiah* per month. The underlying food expenditure items are recorded for the week prior to enumeration and scaled up to the monthly level by the factor $30/7$. The underlying non-food expenditure items are recorded for the year prior to enumeration and scaled down to the monthly level by the factor $1/12$. *Below poverty line* is an indicator for whether or not the household's total expenditures per capita fell below the provincial rural or urban poverty line in the given year. Per capita expenditure quintiles are computed separately within the full national sample and within the 31 provinces in which sample households reside. The 2005 quintiles are calculated including attritors. The expenditure figures are not adjusted for inflation between 2005 and 2006. Since we are estimating an outcome expressed as the (log) difference between 2006 and 2005, the inflation rate is subsumed in the constant in all pure OLS estimates.

Table 2: Propensity score model, $\mathbb{P}(D_h > 0|\mathbf{X}_h)$

Regressor	**Coefficient**	**(Std. error)**
Urban Area	-0.177	(0.114)
HH Head Female	0.617***	(0.114)
Land owned (hectares)	-0.099***	(0.031)
Land owned^2 (hectares)	0.001***	(0.000)
HH ever participate Rice for the Poor	0.961***	(0.085)
# children in school	-0.102	(0.075)
# children in school^2	0.023	(0.020)
Indicators for HH size$\in \{2,\ldots,12\}$	—	[0.065]†
Floor area	-0.005***	(0.002)
Household composition (reference=Share Adult Males, 10+ yrs)		
Share Female Children, 0-9 yrs	0.619**	(0.265)
Share Male Children, 0-9 yrs	0.421*	(0.234)
Share Adult Females, 10+ yrs	-0.025	(0.186)
Primary HH income source (reference=other)		
Trade/Retail	-0.179	(0.117)
Financial/Real Estate	-0.782*	(0.428)
Agriculture	0.060	(0.126)
Mining	-0.235	(0.156)
Manufacturing	0.158	(0.125)
Electricity/Gas/Water	0.269	(0.882)
Construction	0.260*	(0.141)
HH head education level (reference=no education)		
Primary	-0.283**	(0.114)
Junior secondary	-0.571***	(0.142)
Senior secondary	-1.091***	(0.147)
Higher	-2.384***	(0.347)
Housing status (reference=other)		
Own house	-0.085	(0.127)
Lease house	-0.132	(0.238)
Rent house	-0.386	(0.258)
Free house	0.015	(0.227)
Official house	-0.719	(0.528)
Roof type (reference=other)		
Concrete roof	-0.849*	(0.451)
Tile roof	-0.418	(0.328)
Shingle roof	-0.490	(0.437)
Iron roof	-0.410	(0.320)
Asbestos roof	-0.425	(0.422)
Fiber/Thatch roof	-0.400	(0.378)
Wall type (reference=other)		
Brick wall	-0.337	(0.304)
Wood wall	0.023	(0.292)
Bamboo wall	0.477	(0.307)
Floor type (reference=other)		
Cement/Tile/Plaster floor	0.133	(0.538)
Wood/Reed/Bamboo floor	0.290	(0.544)
Earthen floor	0.797	(0.549)
Source of drinking water (reference=other)		
Bottled water	-0.978**	(0.394)
Pump water	-1.039***	(0.289)
Tap water	-0.427	(0.335)
Protected well water	-0.678**	(0.272)
Unprotected well water	-0.918***	(0.288)
Protected spring water	-0.985***	(0.306)
Unprotected spring water	-0.883***	(0.322)
River water	-0.929***	(0.322)
Rain water	-0.562	(0.379)
Buy drinking water	-0.166	(0.153)
Toilet facilities (reference=other)		
Own toilet	-0.218*	(0.128)
Shared toilet	-0.016	(0.132)
Public toilet	0.011	(0.220)
Source of light (reference=other)		
PLN electricity	0.061	(0.597)
Non-PLN electricity	-0.082	(0.690)
Pump lantern	0.899	(0.631)
Oil lamp	0.648	(0.595)
Toilet disposal location (reference=other)		
Septic tank	-0.269	(0.175)
Pond/Rice field	-0.044	(0.206)
Lake, river, sea	-0.027	(0.150)
Beach	-0.034	(0.167)
Constant	0.382	(0.938)
Pseudo-R^2		0.22

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. Estimated using balanced panel containing 9050 households from *Susenas* 2005 and 2006 Panel. Standard errors are clustered by village. All variables are as reported in February-April 2005. The regression also controls for province fixed effects. *PLN* is the state-run electricity firm.

Table 3: Idiosyncratic vs. Spatial Variation in Staggering

Fixed Effects	Province	District	Subdistrict	Province	District	Subdistrict
$\mathbf{X}_{h,t-1}$ controls	No	No	No	Yes	Yes	Yes
			specification: $Pr(D > 0)$			
	(1)	(2)	(3)	(4)	(5)	(6)
$H_0: \boldsymbol{\beta}_{\mathbf{X}} = \mathbf{0}$						
F statistic	—	—	—	30.72	28.67	28.38
$[p - value]$	—	—	—	$[< 0.001]$	$[< 0.001]$	$[< 0.001]$
R^2	0.049	0.167	0.237	0.236	0.326	0.385
			specification: $Pr(D = 2 \mid D > 0)$			
	(7)	(8)	(9)	(10)	(11)	(12)
$H_0: \boldsymbol{\beta}_{\mathbf{X}} = \mathbf{0}$						
F statistic	—	—	—	2.93	1.66	0.90
$[p - value]$	—	—	—	$[< 0.001]$	$[< 0.001]$	$[0.705]$
R^2	0.264	0.811	0.893	0.325	0.821	0.896

Notes: D denotes the number of disbursements received by *Susenas* enumeration in early 2006. Linear probability regressions for $Pr(D = 2|\cdot)$ and $Pr(D > 0|\cdot)$ are based on the sample of recipient and all households, respectively. There are 30 provinces, 339 districts, and 619 subdistricts.

Table 4: Staggering is Orthogonal to Interregional Differences

	(1)	(2)	(3)	(4)	(5)	(6)
log distance to subdistrict capital	0.042	0.027	0.032	0.039	0.039	0.042
	(0.037)	(0.039)	(0.039)	(0.039)	(0.039)	(0.039)
log distance to district capital	-0.057	-0.053	-0.040	-0.033	-0.034	-0.033
	(0.032)*	(0.033)	(0.036)	(0.036)	(0.036)	(0.036)
log distance to Jakarta	0.005	-0.034	-0.033	-0.030	-0.031	-0.022
	(0.032)	(0.041)	(0.041)	(0.041)	(0.040)	(0.043)
log district population, 2005		-0.072	-0.075	-0.076	-0.076	-0.082
		(0.046)	(0.046)	(0.045)*	(0.045)*	(0.046)*
urban village			0.064	0.067	0.069	0.074
			(0.050)	(0.049)	(0.051)	(0.051)
village road paved				0.024	0.025	0.026
				(0.054)	(0.054)	(0.054)
village accessible only by water				-0.100	-0.100	-0.107
				(0.081)	(0.081)	(0.082)
log mean household exp./capita in district, 2005					-0.008	-0.017
					(0.093)	(0.094)
rainfall shock, 2005						0.120
						(0.169)
Number of Households	2,383	2,383	2,383	2,383	2,383	2,383
R^2	0.010	0.022	0.025	0.031	0.031	0.031

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. Linear probability regressions based on the sample of recipient households using the following specification: $Pr(D_{hv} = 2 \mid D_{hv} > 0) = \boldsymbol{\gamma}\mathbf{Z}_v + v_{hv}$, where $\mathbf{Z}_v$ comprises a vector of characteristics associated with the village or region within which household v resides. Distance to (sub)district capital is based on travel distance; distance to Jakarta is great-circle. Standard errors are clustered at the district level in all specifications. The sample decline is due to a loss of villages in Papua for which I could obtain reliable matches. All other results robust to dropping these households.

Table 5: Baseline Estimates of Multi-valued Treatment Effects, Short- and Medium-Term

Estimator	OLS	IPW	Double Robust		Control Function
			$(\hat{P}_h)$	$(\mathbf{X}_h)$	
	(1)	(2)	(3)	(4)	(5)
	Short-Term: 2005-2006				
τ_{10}: receipt of disbursement 1	-0.064	-0.089	-0.089	-0.089	-0.075
	(0.027)**	(0.035)**	(0.034)***	(0.030)***	(0.030)**
τ_{21}: receipt of disbursement 2	0.051	0.073	0.075	0.070	0.076
	(0.030)*	(0.036)**	(0.035)**	(0.032)**	(0.033)**
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	-0.013	-0.016	-0.014	-0.019	0.001
	(0.014)	(0.021)	(0.020)	(0.017)	(0.017)
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	No	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No
Number of Households	9,010	9,010	9,010	9,010	9,010
R^2	0.045	0.088	0.091	0.170	0.104
	Medium-Term: 2005-2007				
τ_{10}: receipt of disbursement 1	-0.037	-0.057	-0.066	-0.045	-0.025
	(0.040)	(0.039)	(0.039)*	(0.034)	(0.038)
τ_{21}: receipt of disbursement 2	0.029	0.032	0.035	0.009	0.031
	(0.045)	(0.044)	(0.043)	(0.038)	(0.042)
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	-0.008	-0.026	-0.031	-0.036	0.006
	(0.020)	(0.024)	(0.024)	(0.022)	(0.022)
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	No	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No
Number of Households	6,992	6,992	6,992	6,992	6,992
R^2	0.044	0.056	0.062	0.146	0.069

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. The dependent variable in all specifications is Δlog total household expenditures per capita between 2005 and 2006/2007. Columns 2-5 are estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$, where the normalization is over the entire sample for the given time horizon. Column 3 controls linearly for the propensity score and column 5 for a fifth-order polynomial in the propensity score allowing it to vary by treatment and control. Column 4 controls for all covariates $\mathbf{X}_h$ used to estimate the propensity score. Standard errors clustered by village. All columns include province fixed effects.

Table 6: Multi-valued Treatment Effects by Expenditure Type

Growth Horizon⟶	2005-2006			2005-2007		
Expenditure Type⟶	**total**	**food**	**non-food**	**total**	**food**	**non-food**
	(1)	(2)	(3)	(4)	(5)	(6)
τ_{10}: receipt of disbursement 1	-0.075	-0.093	-0.048	-0.025	-0.029	0.008
	(0.030)**	(0.030)***	(0.047)	(0.038)	(0.035)	(0.052)
τ_{21}: receipt of disbursement 2	0.076	0.097	0.036	0.031	0.068	-0.033
	(0.033)**	(0.034)***	(0.050)	(0.042)	(0.040)*	(0.057)
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	0.001	0.005	-0.013	0.006	0.039	-0.024
	(0.017)	(0.017)	(0.024)	(0.022)	(0.021)*	(0.033)
Reweighted	Yes	Yes	Yes	Yes	Yes	Yes
Propensity Score Polynomial	Yes	Yes	Yes	Yes	Yes	Yes
Number of Households	9,010	9,010	9008	6,992	6,992	6,992

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. The dependent variable in all specifications is Δlog household expenditures on the given commodity group per capita between 2005 and 2006/2007. All columns estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$. All columns include a 5th order polynomial in the propensity scores that is allowed to vary by treatment and control. Standard errors clustered by village. All columns include province fixed effects.

Table 7: Multi-valued Treatment Effects by Disaggregated Expenditure Group

	τ_{10}	τ_{21}	τ_{20}	No. of Households
total	-0.075 (0.030)**	0.076 (0.033)**	0.001 (0.017)	9,010
food	-0.093 (0.030)***	0.097 (0.034)***	0.005 (0.017)	9,010
rice	-0.039 (0.038)	0.056 (0.044)	0.017 (0.023)	8,777
tubers	0.075 (0.075)	0.005 (0.086)	0.079 (0.055)	2,733
fish, meat, dairy	-0.18 (0.063)***	0.125 (0.070)*	-0.055 (0.034)	8,338
fruit, nuts, vegetables	-0.128 (0.046)***	0.121 (0.051)**	-0.007 (0.021)	8,850
other	-0.083 (0.039)**	0.088 (0.045)*	-0.005 (0.021)	8,885
prepared food	-0.226 (0.070)***	0.154 (0.078)**	-0.072 (0.037)*	7,653
alcohol, tobacco	-0.185 (0.069)***	0.155 (0.081)*	-0.031 (0.049)	5,330
nonfood	-0.077 (0.058)	0.042 (0.063)	-0.035 (0.030)	9,008
education, health	-0.158 (0.112)	0.189 (0.128)	0.031 (0.063)	6,507
housing, utilities	0.010 (0.054)	-0.030 (0.057)	-.019 (0.024)	9,008
transport, communication	-0.306 (0.105)***	0.269 (0.117)**	-0.033 (0.056)	5,480
appliances	-0.014 (0.067)	0.126 (0.077)*	0.112 (0.044)**	8,904
debt, taxes	0.065 (0.124)	-0.145 (0.137)	-0.079 (0.056)	5,997

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. Each row corresponds to a separate regression with the log difference in the given expenditure category on the left hand side. All rows estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$. All estimates include a 5th order polynomial in the propensity scores that is allowed to vary by treatment and control. Standard errors clustered by village. All estimates include province fixed effects.

Table 8: UCT Benefits Had No Effect on Household Size

Estimator	OLS	IPW	Double Robust		Control Function
			$(\hat{P}_h)$	$(\mathbf{X}_h)$	
	(1)	(2)	(3)	(4)	(5)
τ_{10}: receipt of disbursement 1	0.004	0.033	0.033	0.013	0.031
	(0.057)	(0.065)	(0.065)	(0.066)	(0.062)
τ_{21}: receipt of disbursement 2	-0.006	-0.036	-0.036	-0.011	-0.050
	(0.063)	(0.067)	(0.068)	(0.069)	(0.068)
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	No	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No
Number of Households	9,010	9,010	9,010	9,010	9,010
R^2	0.005	0.011	0.011	0.054	0.015

Notes: All columns estimated by linear probability regressions with Δ log household size between 2005 and 2006 on the left hand side. Columns 2-5 are estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$. Standard errors clustered by village. All columns include province fixed effects.

Table 9: Idiosyncratic vs. Spatial Variation in the "Tax" on UCT Recipients

	(1)	(2)	(3)	(4)
household size, $t-1$	-0.001	0.001	0.003	-0.000
	(0.004)	(0.002)	(0.005)	(0.002)
Fixed Effects (FE)	Province	Subdistrict	Province	Subdistrict
$\mathbf{X}_{h,t-1}$ controls	No	No	Yes	Yes
p-value joint statistical significance			[0.52]	[0.99]
Number of Households	2,410	2,410	2,410	2,410
R^2	0.113	0.822	0.187	0.827

Notes: All columns estimated by linear probability regressions of the following specification: $Pr(transfer_h < full\ amount \mid D) = \boldsymbol{\beta}\mathbf{X}_{h,t-1} + \theta_{FE} + e_h$, where $\mathbf{X}_{h,t-1}$ includes all the baseline household characteristics used to estimate propensity scores. Standard errors clustered by village.

Table 10: Intensive Margin Treatment Effects by Expenditure Group

	(1)	(2)	(3)	(4)
	Dep. Var.: Δ*log total expenditures/capita*			
transfers per capita (000,000s Rp)	0.045	0.045	0.038	0.066
	(0.008)***	(0.008)***	(0.008)***	(0.011)***
	Dep. Var.: Δ*log food expenditures/capita*			
transfers per capita (000,000s Rp)	0.045	0.045	0.040	0.066
	(0.008)***	(0.009)***	(0.009)***	(0.011)***
	Dep. Var.: Δ*log non-food expenditures/capita*			
transfers per capita (000,000s Rp)	0.056	0.056	0.049	0.091
	(0.013)***	(0.013)***	(0.015)***	(0.023)***
Treatment Indicators	Yes	Yes	Yes	Yes
Reweighted	No	Yes	Yes	Yes
Propensity Score Polynomial	No	No	Yes	Yes
Household Size Indicators	No	No	No	Yes
Number of Households	9,010	9,010	9,010	9,010
R^2	0.106	0.121	0.106	0.121

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. Each cell corresponds to a separate regression. Transfers are rescaled to 100,000s of Rupiah (approx. 10 USD). Columns 2-4 are estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$. Standard errors clustered by village. All columns include province fixed effects.

Table 11: Poverty Transitions and the UCT (Multinomial Logit AME)

	(1) chronic poor $Pr(poor_{t-1}=1, poor_t=1)$	(2) into poverty $Pr(poor_{t-1}=0, poor_t=1)$	(3) out of poverty $Pr(poor_{t-1}=1, poor_t=0)$	(4) never poor $Pr(poor_{t-1}=0, poor_t=0)$
	Short-Term: 2005-2006			
τ_{10}: receipt of disbursement 1	0.218	0.099	0.014	-0.331
	(0.047)***	(0.023)***	(0.021)	(0.041)***
τ_{21}: receipt of disbursement 2	0.003	-0.023	0.040	-0.020
	(0.026)	(0.019)	(0.022)*	(0.033)
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	0.221	0.076	0.053	-0.351
	(0.042)***	(0.076)***	(0.053)***	(0.033)***
Reweighted	Yes	Yes	Yes	Yes
Propensity Score Polynomial	Yes	Yes	Yes	Yes
Actual Probability	0.081	0.084	0.063	0.772
Predict Probability	0.080	0.085	0.065	0.770
Number of Households	9,010	9,010	9,010	9,010
	Medium-Term: 2005-2007			
τ_{10}: receipt of disbursement 1	0.086	0.112	0.105	-0.303
	(0.021)***	(0.030)***	(0.029)***	(0.038)***
τ_{21}: receipt of disbursement 2	0.014	0.009	0.024	-0.047
	(0.017)	(0.017)	(0.027)	(0.034)
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	0.099	0.121	0.129	-0.349
	(0.017)***	(0.028)***	(0.018)***	(0.028)***
Reweighted	Yes	Yes	Yes	Yes
Propensity Score Polynomial	Yes	Yes	Yes	Yes
Actual Probability	0.034	0.054	0.110	0.803
Predicted Probability	0.040	0.066	0.107	0.786
Number of Households	6,992	6,992	6,992	6,992

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. The poverty line varies across district $\times$ urban or rural administrative divisions. The average marginal effects (AME) are based on multinomial logit (base outcome is "into poverty") where the weights for treatment households equal one and the weights for control households are given by the normalized $w(\cdot) = \hat{P}_h/(1-\hat{P}_h)$. Standard errors clustered by village. The regression includes province fixed effects.

Table 12: Multi-valued Treatment Effects on Labor Supply, 2005-6

Estimator	OLS	IPW	Double Robust		Control
			$(\hat{P}_h)$	$(\mathbf{X}_h)$	Function
	(1)	(2)	(3)	(4)	(5)
			Short-Term: 2005-2006		
τ_{10}: receipt of disbursement 1	-0.535	-1.710	-1.700	-1.838	-0.437
	(0.737)	(0.887)*	(0.859)**	(0.842)**	(0.773)
τ_{21}: receipt of disbursement 2	0.422	1.168	1.240	1.411	0.843
	(0.846)	(0.934)	(0.927)	(0.920)	(0.874)
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	-0.113	-0.543	-0.461	-0.427	0.406
	(0.451)	(0.585)	(0.563)	(0.532)	(0.525)
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	No	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No
Number of Households	9,010	9,010	9,010	9,010	9,010
R^2	0.015	0.027	0.031	0.102	0.051
			Medium-Term: 2005-2007		
τ_{10}: receipt of disbursement 1	-2.309	-2.571	-2.623	-2.285	-2.565
	(0.946)**	(1.085)**	(1.100)**	(1.105)**	(1.053)**
τ_{21}: receipt of disbursement 2	2.061	2.129	2.148	2.111	2.114
	(1.065)*	(1.167)*	(1.163)*	(1.157)*	(1.179)*
$\tau_{20} \equiv \tau_{21} - \tau_{10}$	-0.248	-0.442	-0.475	-0.173	-0.451
	(0.577)	(0.669)	(0.686)	(0.666)	(0.599)
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	Yes	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No
Number of Households	6,992	6,992	6,992	6,992	6,992
R^2	0.013	0.017	0.018	0.074	0.021

Notes: Significance levels: $*$ 10% $**$ 5% $***$ 1%. The dependent variable in all specifications is Δweekly hours worked per adult between 2005 and 2006, which is calculated as total hours worked divided by number of adult household members. Columns 2-5 are estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1 - \hat{P}_h)$. Standard errors clustered by village. All columns include province fixed effects.

Table 13: Intensive Margin Treatment Effects on Labor Supply

Estimator	OLS	IPW	Double Robust		Control Function
			$(\hat{P}_h)$	$(\mathbf{X}_h)$	
	(1)	(2)	(3)	(4)	(5)
			Δ 2005-2006		
transfers per capita (000,000s Rp)	0.373	-0.592	-0.609	-0.391	-0.202
	(0.363)	(0.610)	(0.603)	(0.481)	(0.458)
			Δ 2005-2007		
transfers per capita (000,000s Rp)	-0.329	-0.406	-0.437	-0.256	-0.280
	(0.483)	(0.561)	(0.566)	(0.531)	(0.561)
Treatment Indicators	Yes	Yes	Yes	Yes	Yes
Household Size Indicators	Yes	Yes	Yes	Yes	Yes
Reweighted	No	Yes	Yes	Yes	Yes
Propensity Score Control(s)	No	Yes	Yes	No	Yes
$\mathbf{X}_h$ Controls	No	No	No	Yes	No

Notes: Significance levels: * 10% ** 5% *** 1%. The dependent variable in all specifications is Δweekly hours worked per adult between 2005 and 2007, which is calculated as total hours worked divided by number of adult household members. Columns 2-5 are estimated by weighted least squares where the weights for treatment households equal one and the weights for control households are given by the normalized $\omega = \hat{P}_h/(1-\hat{P}_h)$. Standard errors clustered by village. All columns include province fixed effects.

Table 14: (Agricultural Household) Expenditures Respond to Transitory Rainfall Shocks

	(1)	(2)	(3)	(4)	(5)
log rainfall mean deviation	0.031	-0.048	-0.041	-0.055	-0.028
	(0.074)	(0.100)	(0.078)	(0.080)	(0.072)
1(rural village)		-0.005			
		(0.020)			
1(rural village) × rainfall shock		0.126			
		(0.112)			
1(agriculture primary income)			0.017		
			(0.015)		
1(agri. primary income) × rainfall shock			0.218		
			(0.090)**		
1(own any agri. land)				-0.001	
				(0.015)	
1(own any agri. land) × rainfall shock				0.171	
				(0.094)*	
agri. land (Ha)					0.001
					(0.003)
agri. land (Ha) × rainfall shock					0.063
					(0.029)**
Number of Households	8,922	8,922	8,922	8,922	8,922
R^2	0.042	0.043	0.044	0.043	0.045

Notes: Significance levels: * 10% ** 5% *** 1%. The dependent variable in all specifications is Δlog total household expenditures per capita between 2005 and 2006. The rainfall shock is the log deviation of the seasonal rainfall level in the district from the long-run (1952-2004) district mean. Standard errors clustered by district. All columns include province fixed effects. The interaction terms are as observed at baseline.

Appendix

A Propensity Scores and Reconstructed Quasi-PMT Scores

To estimate the probability that household h receives treatment d, $\mathbb{P}(D_h = d \mid \mathbf{X}_h)$, we consider the following specification, which roughly approximates information on household h available to enumerators and local officials in mid-2005,

$$\mathbb{P}(D_h = d) = F\left(\boldsymbol{\beta}\mathbf{X}_h^{\text{fam}} + \boldsymbol{\gamma}\mathbf{X}_h^{\text{house}} + \boldsymbol{\alpha}\mathbf{X}_h^{\text{head}} + \boldsymbol{\delta}\mathbf{X}_h^{\text{welfare}} + \zeta_h^d > 0\right), \tag{11}$$

All right-hand variables are observed in February 2005: $\mathbf{X}_h^{\text{fam}}$ is a vector of demographic variables including household age structure, gender breakdown; $\mathbf{X}_h^{\text{house}}$ contains variables pertaining to the quality of the physical structures in which household h lives; $\mathbf{X}_h^{\text{head}}$ are characteristics of the head of the household, $\mathbf{X}_h^{\text{welfare}}$ contain indicators for employment among household members, prior participation in government welfare programs, and amount of land owned; F is the relevant CDF; and ζ_h captures all variables unobservable to the econometrician but possibly observable to program administrators. We also control for province fixed effects to subsume some of the regional differences in targeting infrastructure (among other things). A full elaboration of the coefficient estimates was reported in Table 2.[29] Given our large set of dummy variables, there is little advantage estimating equation (11) nonparametrically.[30]

As discussed in Section 3.1, although we made every effort to reconstruct the underlying PMT scores using available data, the resulting scores were not discriminating enough to allow for even a fuzzy regression-discontinuity research design. After transforming applicable questions in *Susenas* 2005 into the corresponding variable-specific eligibility criteria, we apply the district-specific PMT coefficients corresponding to the given variables to produce a measure $\widetilde{\mathbb{P}}_h$. This variable reflects a data-constrained approximation to the actual PMT scores based on the original eligibility survey.[31] According to program guidelines, households with PMT scores above the 70th percentile should qualify for benefits. We take this rule to our estimates $\widetilde{\mathbb{P}}_h$ in search of a potential discontinuity. Unfortunately, as seen in Figure A7, no such discontinuity can be found—perhaps unsurprisingly given the evidence on leakage and undercoverage. Moreover, the actual probability of UCT receipt looks quite similar across the distribution of the estimated propensity scores $\hat{\mathbb{P}}_h$. Yet, if we predict the probability of program receipt using $\widetilde{\mathbb{P}}_h$ as the only regressor—effectively fixing $(\beta, \gamma, \delta, \alpha)$ in equation 11 at the district-specific PMT coefficients—and accordingly reweight households in the control group, the balance at baseline is much worse than when using our arguably more flexible approach based on a richer set of variables plausibly in the information set of local officials engaged in community-based alongside or possibly in defiance of official targeting. This can be seen by comparing the effect of reweighting the control group in Figure A8, which uses $\widetilde{\mathbb{P}}_h$, and Figure 5 discussed in the paper, which uses our estimated propensity scores.

[29] The official eligibility survey grouped several response categories to questions in *Susenas* concerning household characteristics. Whether one leaves the individual responses as separate indicators (in a fully saturated sense) or groups them according to the rubric in the original survey does not matter for the qualitative findings presented below.

[30] Doing so using the Klein and Spady (1993) estimator yields an estimated propensity score that has a 0.95 correlation with the simpler parametric logit.

[31] Prior to this, we rescale the coefficients to ensure that they sum to 1 after dropping the questions not available in *Susenas*.

Figure A7: **Comparing Propensity Score Estimates and Approximated Quasi-PMT Scores**

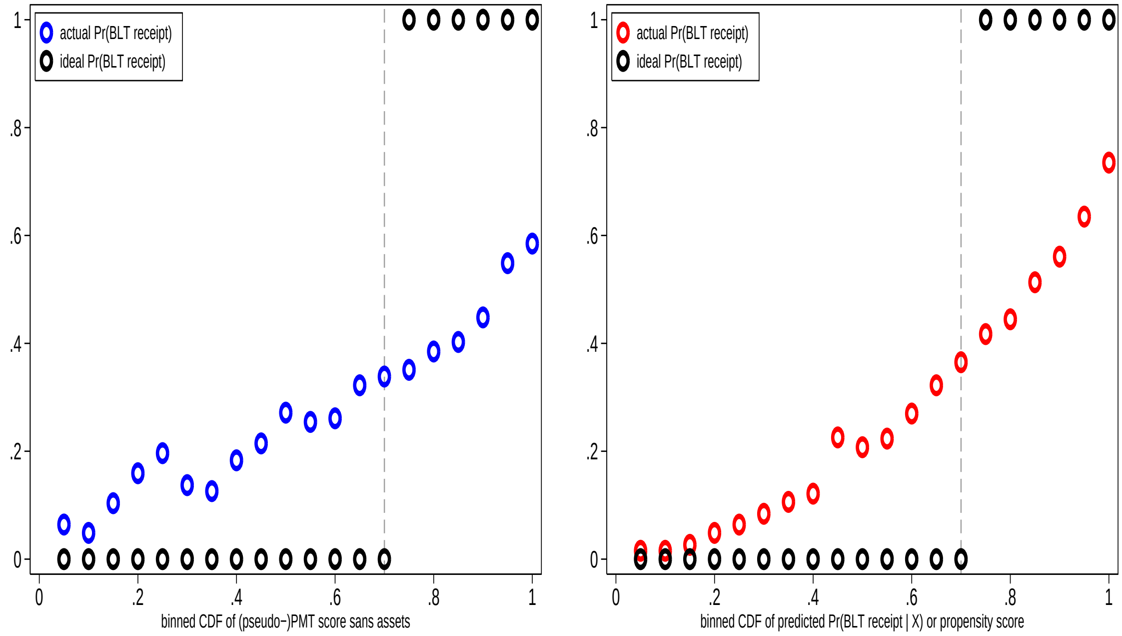

Notes: LEFT—The circles capture the share of UCT (BLT) recipients within the given bin where the bins are 0.05 width slices of the CDF of the quasi-PMT scores approximated using the procedures described in the text. The dashed vertical line constitutes the 30% threshold above which households were (in theory) supposed to receive the program. RIGHT—The circles capture the share of UCT (BLT) recipients within the given bin where the bins are 0.05 width slices of the CDF of the propensity scores obtained from estimating a binary version of equation (11) by maximum likelihood where ζ_h is logistic distributed.

Figure A8: Baseline Expenditure Distributions by Treatment Status

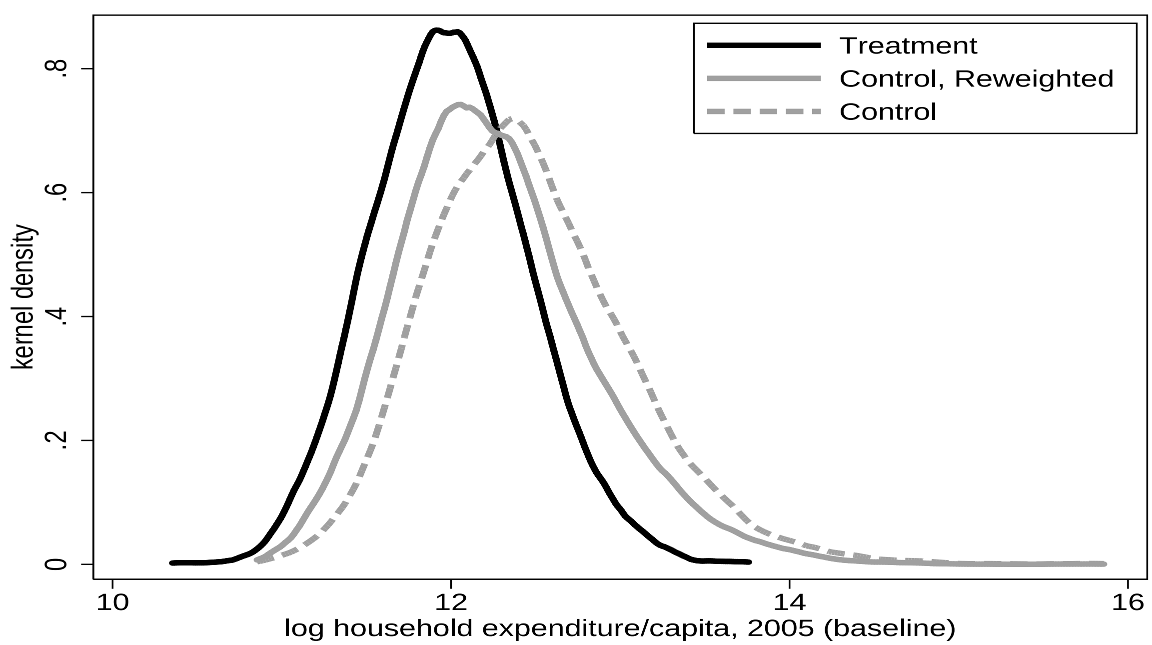

Notes: All distributions estimated using Epanechnikov kernel and a rule-of-thumb bandwidth. The "Control (Reweighted)" observations are adjusted using inverse probability weights (IPW) based on normalized estimated odds of treatment $\omega = \widetilde{P}/(1-\widetilde{P})$, where $\widetilde{\mathbb{P}}_h$ is as described in Appendix A.