Targeting the Poor: Evidence from a Field Experiment in Indonesia[†]

By Vivi Alatas, Abhijit Banerjee, Rema Hanna,
Benjamin A. Olken, and Julia Tobias*

This paper reports an experiment in 640 Indonesian villages on three approaches to target the poor: proxy means tests (PMT), where assets are used to predict consumption; community targeting, where villagers rank everyone from richest to poorest; and a hybrid. Defining poverty based on PPP$2 per capita consumption, community targeting and the hybrid perform somewhat worse in identifying the poor than PMT, though not by enough to significantly affect poverty outcomes for a typical program. Elite capture does not explain these results. Instead, communities appear to apply a different concept of poverty. Consistent with this finding, community targeting results in higher satisfaction. (JEL C93, I32, I38, O12, O15, O18, R23)

Targeted social safety net programs have become an increasingly common tool to address poverty (Coady, Grosh, and Hoddinott 2004). In developed countries, the selection of the beneficiaries for these programs ("targeting") is frequently accomplished through means-testing: only those with incomes below a certain threshold are eligible. In developing countries, however, where most potential recipients work in the informal sector and lack verifiable records of their earnings, credibly implementing a conventional means test is challenging.

Consequently, in developing countries, there is an increased emphasis on targeting strategies that do not rely on directly observing incomes. In particular, there are two main types of strategies that we consider in this paper: proxy means tests

(PMTs) and community-based targeting.[1] In a PMT, which has been used in the Mexican Progresa/Oportunidades and Colombian Familias en Acción programs, the government collects information on assets and demographic characteristics to create a "proxy" for household consumption or income, and this proxy is in turn used for targeting. In community-based methods, such as the Bangladesh Food-for-Education program (Galasso and Ravallion 2005) and the Albanian Economic Support safety net program (Alderman 2002), the government allows the community or some part of it (e.g., local leaders) to select the beneficiaries. Both methods aim to address the problem of unobservable incomes. In the PMTs, the presumption is that household assets are harder to conceal from government surveyors than income; in community-based targeting, the presumption is that wealth is harder to hide from one's neighbors than from the government.

The choice between the two approaches is generally framed as a trade-off between the better information that communities might have versus the risk of elite capture in the community process. By focusing on assets, PMTs capture the permanent component of consumption. In the process, however, they miss out on transitory or recent shocks. For example, a family may fall into poverty because one of its members has fallen ill and cannot work, but because the family has a large house, a PMT may still classify it as nonpoor. Neighbors, on the other hand, may know the family's true situation by regularly observing the way that they live.[2] If the community perceives that the PMT is wrong, a lack of legitimacy and political instability may ensue.[3]

While community targeting allows for the use of better local information, however, it also opens up the possibility that targeting decisions may be based on factors beyond poverty as defined by the government. This may be due to genuine disagreements about what "poverty" means: the central government typically evaluates households based on consumption, whereas the utility function used by local communities may include other factors, such as a household's earning potential, nonincome dimensions of poverty, or its number of dependents.[4] Or, the two groups may place a different weight on the same variable when predicting consumption. Moreover, the community process could also favor the friends and relatives of the elites, resulting in a lack of legitimacy with the process.

Given the trade-offs involved, which method works best is ultimately an empirical question. If elite capture of community targeting is important, then the PMT could dominate community targeting either based on the government's consumption-based metric or a more holistic welfare metric, since the PMT limits the

---

[1] Self-targeting, where individuals self-identify as poor and then are subject to verification (as in Nichols and Zeckhauser 1982) is also increasingly being used in the developing world. While we are unable to address self-targeting techniques in this paper, this is the focus of our future work.

[2] Seabright (1996) makes the theoretical argument that greater local information is one of the advantages of the community methods. Alderman (2002) and Galasso and Ravallion (2005) provide empirical evidence that communities may have additional information beyond the PMTs.

[3] See, for example, "Data Penerima BLT di Semarang Membingungkan" (BLT Beneficiary List in Semarang Confuses), *Kompas*, May 15, 2008; "Old data disrupts cash aid delivery," *Jakarta Post*, September 6, 2008; "Poorest still waiting for cash aid," *Jakarta Post*, June 24, 2008; "Thousands protest fuel plan, cash assistance," *Jakarta Post* May 22, 2008.

[4] There is little existing evidence that what we perceive as "targeting errors" may be due to different conceptions of poverty by the different stakeholders involved. One exception is Ravallion (2008), which shows that the objective function of the program administrators for a targeted welfare program in China held a broader concept of poverty than that of economists/evaluators.

opportunity for capture. If better local information is important, then community targeting could dominate the PMT on both of these metrics. If a different local conception of welfare is empirically important, then the PMT may best match the government's consumption-based metric, while community targeting may work best based on alternative welfare metrics.[5] In this paper, we use randomized evaluation techniques to compare PMT targeting with methodologies that allow for varying degrees of community inclusion in the decision-making process. We first compare how the methods perform from the perspective of the central government: poverty as measured by per capita expenditures and satisfaction with the targeting process. To understand why the methods produce different results, we then investigate the trade-offs discussed above along four dimensions: elite capture, the role of effort, differences in information, and different conceptions of poverty.

In 640 villages in Indonesia, we conducted a field experiment in collaboration with the government. The government, through the Central Bureau of Statistics, implemented a cash transfer program that sought to distribute Rp. 30,000 (about $3) to households that fell below location-specific poverty lines. In a randomly selected one-third of the villages, the government conducted a PMT in order to identify the beneficiaries. In another third of these villages, chosen at random, it employed community targeting (hereafter, the community method): the community members were asked to rank everyone from richest to poorest during a meeting, and this ranking determined eligibility. In the remaining villages, it used a combination of the two methods (hereafter, the hybrid method): communities engaged in the ranking exercise, and then the ranks were used to limit the universe of households whom the government would survey. Eligibility was then determined by conducting PMT on this limited list. This hybrid aimed to utilize the communities' knowledge, while using the PMT as a check on potential elite capture.

We begin by evaluating the methods from the perspective of the central government; i.e., which method best targeted the poor based on consumption-based poverty and which method produced the highest satisfaction with the beneficiary list. We conducted a baseline survey that collected per capita expenditure data from a set of households prior to the experiment and then defined a household as poor if it fell below the PPP$2 per day cutoff. We find that both the community and hybrid methods perform worse than the PMT on this metric: in both methods, there was a three percentage point (ten percent) increase in the error rate based on consumption (which we will call "error rate" from now on for conciseness) relative to the PMT. However, if we focus on the very poorest households within the poor category, the community-based strategies actually perform as well (if not better).

On net, the differences in targeting accuracy across the methods are not large; for example, for a typically sized transfer program in Indonesia, simulations suggest that these different targeting methods would not yield significantly different effects

---

[5] Coady, Grosh, and Hoddinott (2004) conduct a meta-analysis of 111 targeted antipoverty programs, including 7 PMTs and 14 cases of community-based targeting. They find no difference in the performance of these two models, as measured by the fraction of resources that went to the bottom 40 percent. As the authors point out, however, two sources of bias complicate the interpretation of these results. First, community targeting is often chosen when state capacity is limited. In such places, the PMT would have fared worse had it been tried. Second, many small projects have used the community model, but fail to systematically report data. Thus, the included examples of community-based targeting tend to be bigger and, potentially, better run.

on reducing the poverty rate in Indonesia. Finally, we find that the results are similar in both urban and rural locations, in villages with greater or less inequality, and with greater or lesser levels of social connectedness; this suggests that the results may be generalizable along these dimensions.

Despite the somewhat worse targeting outcomes based on consumption, the community methods resulted in higher satisfaction and greater legitimacy of the process along all of the dimensions that we considered. Community targeting resulted in 60 percent fewer complaints than the PMT, and there were many fewer difficulties in distributing the actual funds in the community treatment villages. When asked ex post about the targeting results, villagers in community treatment villages suggested fewer modifications to the beneficiary list.

We next turn to understanding why the community methods may differ from the PMT. We consider four dimensions: elite capture, community effort, local concepts of poverty, and information. To test for elite capture in the community treatment, we randomly divided the community and hybrid villages so that, in half of these villages, everyone in the community was invited to participate in the ranking meeting, whereas in the other half, only the "elites" (i.e., local community leaders such as the subvillage head, teachers, religious leaders, etc.) were invited. In addition, we gathered data in the baseline survey on which households were related to the local elites. We find no evidence of elite capture. The error rates were the same, regardless of whether only the elites attended the meeting. Moreover, we find no evidence that households that are related to the elites are more likely to receive funds in the community treatments relative to the PMT. In fact, we find the opposite: in the community treatments, elites and their relatives are much less likely to be put on the beneficiary list, regardless of their actual income levels.

To examine the role of effort, we randomized the order in which households were considered at the meetings. This allows us to test whether the effectiveness of community targeting differs between households that were ranked first and those ranked last (when fatigue may have set in). We find that effort matters: at the start of the community meeting, targeting performance is better than in the PMT, but it worsens as the meeting proceeds.

To examine the role of preferences and information, we studied alternative metrics of evaluating perceptions of poverty from our baseline survey. First, we asked every survey respondent to rank a set of randomly chosen villagers from rich to poor (hereafter, "survey ranks"). Second, we asked the head of the subvillage to conduct the same exercise. Finally, and perhaps most importantly, we asked each household that we interviewed to assess its own welfare level subjectively. We find that the community treatment produces a ranking of villagers that is much more correlated with these three alternate metrics than the ranking produced by PMT. In other words, the community treatments moved the targeting outcomes away from a ranking based purely on per capita consumption and toward the rankings that one would obtain by polling different classes of villagers or by asking villagers to rate themselves.

There are two ways of explaining these findings: either the community has less information about different households' per capita consumption than the PMT, or the community's conception of poverty is different from that based solely on per capita consumption. The evidence suggests that the latter theory predominantly drives the results. First, even controlling for all variables in the PMT, the community

members' rankings of other households in the village contain information about those households' per capita consumption, which shows that community members have residual information about consumption beyond that contained in the PMT variables. Second, when we investigate how the survey ranks differ from consumption, we find that communities place greater weight on factors that predict earnings capacity than would be implied by per capita consumption. For example, conditional on actual per capita consumption, the communities consider widowed households poorer than the typical household. The fact that communities employ a different concept of poverty explains why community targeting performance might differ from the PMT, as well as why it results in greater satisfaction levels.

The paper proceeds as follows. We discuss the empirical design in Section I, and describe the data in Section II. In Section III, we compare how each of the main targeting methods fared in identifying the poor. Section IV tests for evidence of elite capture, while Section V aims to understand the role of effort. In Section VI, we test whether the community and the government have different maximands. Section VII explores the differences in the community's maximand in greater depth. Section VIII concludes.

I. Experimental Design

A. *Setting*

This project occurred in Indonesia, which is home to one of the largest targeted cash transfer programs in the developing world, the Direct Cash Assistance (*Bantuan Langsung Tunai* (BLT)) program. Launched in 2005 and renewed again in 2008, the BLT program provided transfers of about US$10 per month to about 19.2 million households during periods of economic crisis. The targeting in this program was accomplished through a combination of community-based methods and PMTs. Specifically, the Central Statistics Bureau (*Badan Pusat Statistik* (BPS)) enumerators met with neighborhood leaders to create a list of households that could potentially qualify for the program. The BPS enumerators then conducted an asset survey and a PMT for the listed households.

Targeting has been identified by policymakers as one of the key problems in the BLT program. Comparing with the goal of targeting the poorest one-third of households, the World Bank estimates that 45 percent of the funds were incorrectly provided to nonpoor households and 47 percent of the poor were excluded from the program in 2005–2006 (World Bank 2006).[6] Perhaps more worrisome from the government's perspective is the fact that citizens voiced substantial dissatisfaction with the beneficiary lists. Protests about mistargeting led some village leaders to resign rather than defend the beneficiary lists to their constituents: over 2,000 village officials refused to participate in the program for this reason.[7] The experiment reported in this paper was

---

[6] Targeting inaccuracy has been documented in many government antipoverty programs (see, for example, Olken 2006; Daly and Fane 2002; Cameron 2002; and Conn et al. 2008).

[7] See, for example, "BLT Bisa Munculkan Konflik Baru" (BLT may create new conflicts), *Kompas,* May 17, 2008; "Kepala Desa Trauma BLT" (A village head's trauma with BLT), *Kompas,* May 24, 2008; "Ribuan Perangkat Desa Tolak Salurkan BLT" (Thousands of village officials refuse to distribute BLT), *Kompas,* May 22, 2008; and "DPRD Indramayu Tolak BLT" (District parliament of Indramayu refuses BLT), *Kompas,* May 24, 2008.

designed and conducted in collaboration with BPS to investigate these two primary targeting issues: targeting performance and popular acceptance of the targeting results.

B. *Sample*

The sample for the experiment consists of 640 subvillages spread across three Indonesian provinces: North Sumatra, South Sulawesi, and Central Java. The provinces were chosen to represent a broad spectrum of Indonesia's diverse geography and ethnic makeup. Within these three provinces, we randomly selected a total of 640 villages, stratifying the sample to consist of approximately 30 percent urban and 70 percent rural locations.[8] For each village, we obtained a list of the smallest administrative unit within it (a *dusun* in North Sumatra and *Rukun Tetangga* (RT) in South Sulawesi and Central Java), and randomly selected one of these subvillages for the experiment. These subvillage units are best thought of as neighborhoods. Each subvillage contains an average of 54 households and has an elected or appointed administrative head, whom we refer to as the subvillage head.

C. *Experimental Design*

In each subvillage, the Central Statistics Bureau (BPS) and Mitra Samya, an Indonesian NGO, implemented an unconditional cash transfer program, where beneficiary households would receive a one-time, Rp. 30,000 (about $3) cash transfer. The amount of the transfer is equal to about ten percent of the median beneficiary's monthly per capita consumption, or a little more than one day's wage for an average laborer.[9]

Each subvillage was randomly allocated to one of the three targeting treatments that are described in detail below.[10] The number of households that would receive the transfer was set in advance through a geographical targeting approach, such that the fraction of households in a subvillage that would receive the subsidy was held constant, on average, across the treatments. We then observed how each treatment selected the set of beneficiaries.

After the beneficiaries were finalized, the funds were distributed. To publicize the lists, the program staff posted two copies of it in visible locations such as roadside food stalls, mosques, or the subvillage head's house. They also placed a suggestion box and a stack of complaint cards next to the list, along with a reminder about the program details. Depending on the subvillage head's preference, the cash distribution could occur either through door-to-door handouts or at a community meeting. After at least three days, the suggestion box was collected.

---

[8] An additional constraint was applied to the district of Serdang Bedagai because it had particularly large-sized subvillages. All villages in this district with average populations above 100 households per subvillage were excluded. In addition, five of the originally selected villages were replaced prior to the randomization due to an inability to reach households during the baseline survey, the village head's refusal to participate, or conflict.

[9] While the transfer is substantially smaller than in the national BLT program, the amount is nonetheless substantial. For example, in September 2008, more than 20 people were killed during a stampede involving thousands when a local wealthy person offered to give out charity of Rp. 30,000 per person ("21 Orang Tewas demi Rp. 30,000" (21 People Killed by 30,000 Rupiah), *Kompas*, September 15, 2008).

[10] Administrative costs of the 3 programs were $65 per village for the community targeting, $146 for the PMT, and $166 for the hybrid. Including the value of the community members' time, the cost of the community targeting was $110, the cost of the PMT was $153, and the cost of the hybrid was $213.

Main Treatment 1: PMT.—In the PMT treatment, the government created formulas that mapped easily observable household characteristics into a single index using regression techniques. Specifically, it created a list of 49 indicators similar to those used in Indonesia's 2008 registration of poor households, encompassing the household's home attributes (wall type, roof type, etc.), assets (TV, motorbike, etc.), household composition, and household head's education and occupation. Using preexisting survey data, the government estimated the relationship between these variables and household per capita consumption.[11] While it collected the same set of indicators in all regions, the government estimated district-specific formulas due to the high variance in the best predictors of poverty across districts. On average, these regressions had an R2 of 0.48 (online Appendix Table 1).[12]

Government enumerators from BPS collected these indicators from all households in the PMT subvillages by conducting a door-to-door survey. These data were then used to calculate a computer-generated poverty score for each household using the district-specific PMT formula. A list of beneficiaries was generated by selecting the predetermined number of households with the lowest PMT scores in each subvillage.

Main Treatment 2: Community Targeting.—In the community treatment, the subvillage residents determine the list of beneficiaries through a poverty-ranking exercise. To start, a local facilitator visited each subvillage, informed the subvillage head about the program, and set a date for a community meeting. The meeting dates were set several days in advance to allow the facilitator and subvillage head sufficient time to publicize the meeting. Facilitators made door-to-door household visits in order to encourage attendance. On average, 45 percent of households attended the meeting.

At the meeting, the facilitator first explained the program. Next, he displayed a list of all households in the subvillage (from the baseline survey), and asked the attendees to correct the list if necessary. The facilitator then spent 15 minutes helping the community brainstorm a list of characteristics that differentiate the poor households from the wealthy ones in their community.

The facilitator then proceeded with the ranking exercise using a set of randomly ordered index cards that displayed the names of each household in the subvillage. He hung a string from wall to wall, with one end labeled as "most well-off" (*paling mampu*) and the other side labeled as "poorest" (*paling miskin*). Then, he presented the first two name cards from the randomly ordered stack to the community and asked, "Which of these two households is better off?" Based on the community's response, he attached the cards along the string, with the poorer household placed closer to the "poorest" end. Next, he displayed the third card and asked how this household ranked relative to the first two households. The activity continued with the facilitator placing each card one by one on the string

---

[11] Data from Indonesia's SUSENAS (2007) and World Bank's Urban Poverty Project (2007) were used to determine the weights on the PMT formula.

[12] It is possible that a misspecified PMT formula could also generate targeting error. Efforts were made to ensure that indicators were highly predictive of per capita consumption, and the formulas were estimated by districts and urban status to ensure that the weights were appropriate to each area. In addition, it is important to note that the assets and demographic indicators used tend to be similar to indicators used in other settings.

until all the households had been ranked.[13] By and large, the community reached a consensus on the ranks.[14] Before the final ranking was recorded, the facilitator read the ranking aloud so adjustments could be made if necessary.

After all meetings were complete, the facilitators were provided with "beneficiary quotas" for each subvillage based on the geographic targeting procedure. Households ranked below the quota were deemed eligible. Note that prior to the ranking exercise, facilitators told the meeting attendees that the quotas were predetermined by the government, and that all households who were ranked below this quota would receive the transfer. The quota itself was not known by either facilitators or attendees at the time of the meeting. Facilitators also emphasized that the government would not interfere with the community's ranking.

Main Treatment 3: Hybrid.—The hybrid method combines the community ranking procedure with a subsequent PMT verification. In this method, the ranking exercise, described above, was implemented first. There was one key difference, however: at the start of these meetings, the facilitator announced that the lowest-ranked households would be independently checked by government enumerators before the list was finalized. The number of households to be verified was equal to 1.5 times the "beneficiary quota" of households who would later receive the transfer.

After the community meetings were complete, the government enumerators visited the lowest-ranked households to collect the data needed to calculate their PMT scores. Beneficiary lists were then determined using the PMT formulas. Thus, it was possible, for example, that some households could become beneficiaries even if they were ranked as slightly wealthier than the beneficiary quota cutoff line on the community list (and vice versa).

The hybrid treatment aims to take advantage of the relative benefits of both methods. First, as compared to the community method, the hybrid method's additional PMT verification phase may limit elite capture. Second, in the hybrid method, the community is incentivized to accurately rank the poorest households at the bottom of the list, as richer households would later be eliminated by the PMT. Third, as compared to the PMT treatment, the hybrid method's use of the community rankings to narrow the set of households that need to be surveyed may be potentially more cost-effective, in light of the fewer household visits required.

Community Subtreatments.—We designed several subtreatments in order to test three hypotheses about why the results from the community process might differ from those that resulted from the PMT treatment: elite capture, community effort, and within-community heterogeneity in preferences.

First, to test for elite capture, we randomly assigned the community and hybrid subvillages to two groups: a "whole community" subtreatment and an "elite"

---

[13] When at least 10 households had been ranked, the facilitator began comparing each card to the middle card (or, if it was higher than the middle card, to the 75th percentile card), and so on, in order to speed up the process.

[14] If the community did not know a household, or consensus on a household could not be reached, the facilitator and several villagers visited the household after the meeting and added it to the rank list based on the information gained from the visit. In practice, this was done in only 2 of the 431 community or hybrid villages (19 out of 67 households at one meeting, all of whom were boarders at a boarding house, and 5 out of 36 households at the second meeting).

subtreatment. In "whole community" villages, the facilitators actively recruited all community members to participate in the ranking. In the "elite" villages, meeting attendance was restricted to no more than seven invitees that were chosen by the subvillage head. Inviting at least one woman was mandatory and there was some pressure to invite individuals who are usually involved in village decision-making, such as religious leaders or school teachers. The elite meetings are smaller and easier to organize and run. Moreover, the elites may have the legitimacy needed (and possibly even better information) to make difficult choices. The danger of these meetings, however, is that elites might funnel aid to their friends and family (Bardhan and Mookherjee 2005).

Second, we introduced a treatment to test whether the efficacy of the community approach is limited by a community's ability or willingness to expend effort. Specifically, we randomized the order in which households were ranked in order to compare the accuracy at the start and the end of the meeting.[15] The ranking procedure is tedious: on average, it took 1.68 hours. For a subvillage with the mean number of households (54), even an optimal sorting algorithm would require making 6 pairwise comparisons by the time the last card was placed. Thus, by the end of the meetings, the community members may be too tired to rank accurately.

The third set of hypotheses concerns the role of preferences. If the community results differ from the PMT results because of preferences, it is important to understand whether these preferences are broadly shared or are simply a function of who attends the meeting. Meeting times were therefore varied in order to attract different subsets of the community. Half of the meetings were randomly assigned to occur after 7:30 p.m., when men who work during the day could easily attend. The rest were in the afternoon, when we expected higher female attendance. In addition, some meetings were conducted that put a particular focus on "poverty": in half of the meetings, the facilitator led an exercise to identify the ten poorest households in the subvillage before the ranking exercise began (hereafter, "10 poorest treatment").

Randomization Design and Timing.—We randomly assigned each of the 640 subvillages to the treatments as follows (Table 1). In order to ensure experimental balance across the geographic regions, we created 51 geographic strata, where each stratum consists of all villages from one or more subdistricts (*kecamatan*) and is entirely located in a single district (*kabupaten*).[16] Then, we randomly allocated subvillages to one of the three main treatments (PMT, community, or hybrid), stratifying such that the proportion allocated to each was identical (up to integer constraints) within each stratum. We then randomly and independently allocated each community or hybrid subvillage to the subtreatments, stratified both by stratum and main treatment.

From November to December 2008, an independent survey company conducted a census in each subvillage and then collected the baseline data. The targeting treatments and the creation of the beneficiary lists started immediately after the baseline

---

[15] Any new household cards that were added to the stack during this process were ranked last.

[16] Specifically, we first assigned each of the 68 subdistricts (kecamatan) in the sample to a unique stratum. We then took all subdistricts with 5 or fewer sampled subdistricts and merged them with other subdistricts in the same district, so that each of the resulting 51 strata had at least 6 sampled villages.

TABLE 1—RANDOMIZATION DESIGN

Community/hybrid subtreatments			Main treatments		
			Community	Hybrid	PMT
Elite	10 poorest first	Day	24	23	
		Night	26	32	
	No 10 poorest first	Day	29	20	
		Night	29	34	
Whole community	10 poorest first	Day	29	28	
		Night	29	23	
	No 10 poorest first	Day	28	33	
		Night	20	24	
		Total	214	217	209

Notes: This table shows the results of the randomization. Each cell reports the number of sub-villages randomized to each combination of treatments. Note that the randomization of subvillages into main treatments was stratified to be balanced in each of 51 strata. The randomization of community and hybrid subvillages into each subtreatment (elite or full community, 10 poorest prompting or no 10 poorest prompting, and day or night) was conducted independently for each subtreatment, and each randomization was stratified by main treatment and geographic stratum.

survey was completed (December 2008 and January 2009). Fund distribution, the collection of the complaint form boxes, and interviews with the subvillage heads occurred during February 2009. Finally, the survey company conducted the endline survey in late February and early March 2009.

II. Data

A. *Data Collection*

We collected four main sources of data: a baseline household survey, household rankings generated by the treatments, data on the community meeting process (in community/hybrid treatments only), and data on community satisfaction.

Baseline Data.—We conducted a baseline survey in November and December 2008. The survey was administered by SurveyMeter, an independent survey organization. At this point, there was no mention of the experiment to households.[17] We began by constructing a complete list of all households in the subvillage. From this census, we randomly sampled 8 households from each subvillage plus the head of the subvillage, for a total sample size of 5,756 households. To ensure gender balance among survey respondents, in each subvillage, households were randomized as to whether the household head or spouse of the household head would be targeted as the primary respondent. The survey included questions on demographics, family networks in the subvillage, participation in community activities, relationships with local leaders, access to existing social transfer programs, and households' per capita consumption.

---

[17] SurveyMeter enumerators were not told about the targeting experiment.

The baseline survey also included a variety of measures of the household's subjective poverty assessments. In particular, we asked each household to rank the other eight households surveyed in their subvillage from poorest to richest. Finally, we asked respondents several subjective questions to determine how they assessed their own poverty levels.

Data on Treatment Results.—Each of our treatments—PMT, community, and hybrid—produces a rank ordering of all the households in the subvillage (hereafter, the targeting rank list). For the PMT treatment, this is the rank ordering of the PMT score; i.e., predicted per capita expenditures. For the community treatment, it is the rank ordering from the community meetings. For the hybrid treatment, it is the final ranked list (where all households that were verified are ordered based on their PMT score, while those that were not are ordered based on their rank from the community meeting).

Data on Community Meetings.—For the community and hybrid subvillages, we collected data on the meetings' functioning, as well as attendance lists. After each meeting, the facilitators filled out a questionnaire on their perceptions of the community's interest and satisfaction levels.

Data on Community Satisfaction.—After the cash disbursement was complete, we collected data on the community's satisfaction level using four different tools: suggestion boxes, subvillage head interviews, facilitator feedback, and household interviews. First, facilitators placed suggestion boxes in each subvillage along with a stack of complaint cards. Each anonymous complaint card asked three yes/no questions in a simple format: (i) Are you satisfied with the beneficiary list resulting from this program? (ii) Are there any poor households not included on the list? (iii) Are there any nonpoor households included on the list? Second, on the day when the suggestion boxes were collected, the facilitators interviewed the subvillage heads.[18] Third, each facilitator filled out feedback forms on the ease of distributing the transfer payments. Finally, in Central Java province, SurveyMeter conducted an endline survey of three households that were randomly chosen from the eight baseline survey households in each subvillage.

B. *Summary Statistics*

Table 2 provides sample statistics of the key variables. Panel A shows that average monthly per capita expenditures are approximately Rp. 558,000 (about $50). Panel B provides statistics on the errors in targeting based on consumption. By construction, about 30 percent of the households received the cash transfer. We calculated the per capita consumption level in each province (separately by urban and rural areas) that corresponded to the percentage of households that were supposed to receive the

---

[18] We intended to randomly reassign facilitators' designated subvillages after the fund distribution so that no facilitator would collect the subvillage head's feedback from an area that he or she had already visited. While this proved logistically impossible in North Sumatra, the reassignment was implemented in the other provinces.

Table 2—Summary Statistics

Variable	Obs	Mean	SD
Panel A. Consumption from baseline survey			
Per capita consumption (Rp. 1,000s)	5,753	557.501	602.33
Panel B. Mistargeting variables			
On beneficiary list	5,756	0.30	0.46
Error rate based on consumption	5,753	0.32	0.47
Inclusion error (nonpoor = rich + middle)	3,725	0.20	0.40
Exclusion error (poor = near + very poor)	2,028	0.53	0.50
Error rate based on consumption − rich	1,843	0.14	0.35
Error rate based on consumption − middle income	1,882	0.27	0.44
Error rate based on consumption − near poor	1,074	0.59	0.49
Error rate based on consumption − very poor	954	0.46	0.50
Panel C. Rank correlations between treatment results and …			
Per capita consumption	640	0.41	0.34
Community (excluding subvillage head)	640	0.64	0.33
Subvillage head	640	0.58	0.41
Self-assessment	637	0.40	0.34

transfer. This threshold level is approximately equal to the PPP$2 poverty line.[19] We defined the "error rate based on consumption" (hereafter, the error rate) to be equal to 1 if either the household's per capita consumption from the baseline survey was below the threshold line and it did not receive the transfer (exclusion errors) or if it was above the threshold line and did receive it (inclusion errors). We further disaggregate these measures by dividing those below the threshold into the "very poor" and the "near poor," with approximately half of the total poor population in each of these two categories. We likewise divide the population above the threshold in half into the "middle income" and "rich." Based on these metrics, 32 percent of the households were incorrectly targeted based on consumption. Twenty percent of the nonpoor households received it, while 53 percent of the poor were excluded. Reassuringly, errors were less likely to happen for the rich (14 percent), and most likely to happen for the near poor (59 percent).[20]

Panel C provides summary statistics for several alternative metrics that can be used to gauge targeting: the rank correlation for each subvillage between one of four different metrics of household well-being and results of the targeting experiment ("targeting rank list"). This allows us to flexibly examine the relationship between the treatment outcomes and various measures of well-being on a comparable scale. First, we compute the rank correlation with per capita consumption, which tells us how closely the final outcome is to the government's metric of well-being. Second,

---

[19] To see this, note that adjusting the 2005 International Price Comparison Project's PPP-exchange rate for Indonesia for inflation through the end of 2008 yields a PPP exchange rate of PPP$1 = Rp. 5,549 (author's calculations based on World Bank 2008 and the Indonesian CPI). The PPP$2 per-day per-person poverty line therefore corresponds to per capita consumption of Rp. 338,000 per month. In our sample, the average threshold below which households should have received the transfer is Rp. 320,000 per month, or almost exactly PPP$2 per day. The slight discrepancy is due to different regional price deflators used in the geographic targeting procedure.

[20] Measurement error in our consumption survey means that we may overestimate the "true" error rates. Measurement error will be identical in the treatment and control and so it will not affect our estimate of changes in the error rate across treatment conditions.

we compute the rank correlation with the ranks provided by the eight individual households during the baseline survey. This allows us to understand how close the targeting rank list is to the community member's individual beliefs about their fellow community members' well-being. Third, we compute the rank correlation with the ranks provided by the subvillage head in the baseline survey. Finally, we compute the rank correlation with respondents' self-assessment of poverty from the baseline survey.[21] This allows us to understand how closely the treatment result matches individuals' beliefs about their own well-being.

While the targeting rank lists are associated with the consumption rankings, they are more highly associated with the community's rankings of well-being. While the mean rank correlation between the targeting rank lists and the consumption rankings is 0.41, the mean correlation of the targeting rank list with the individual community members' ranks is 0.64, and the correlation with the subvillage head's ranks is 0.58. Finally, we observe a 0.40 correlation between the ranks from the targeted lists with the individuals' self-assessments.

C. *Randomization Balance Check*

To verify that the randomization for the main treatments generates balance across the covariates, we examined the following five characteristics from the baseline survey prior to obtaining the data from the experiment:[22] per capita expenditures, years of education of the household head, calculated PMT score, the share of households that are agricultural, and the years of education of the subvillage head. We also examined five village characteristics from the 2008 PODES, a census of villages conducted by BPS: log number of households, distance to district center in kilometers, log size of the village in hectares, the number of religious buildings per household, and the number of primary schools per household. The results, presented in online Appendix Table 2 and discussed in more detail in the online Appendix, show that the subvillages generally appear to be well balanced.

III. Results on Targeting Performance and Satisfaction

We begin by evaluating the treatments from the government's perspective. Specifically, we examine (i) how the treatments performed in terms of targeting the poor based on per capita consumption; (ii) how the treatments could affect the poverty rate; and (iii) how the treatments performed in terms of satisfaction with and legitimacy of the targeting results.

A. *Targeting Performance Based on Per Capita Consumption*

We begin by comparing how the different targeting methods performed based on per capita consumption levels. Specifically, we compute location-specific pov-

---

[21] Each household was asked, "Please imagine a six-step ladder where on the bottom (the first step) stand the poorest people and on the highest step (the sixth step) stand the richest people. On which step are you today?"

[22] We specified and documented all of the main regressions before examining the data (April 3, 2009); this is available from the authors upon request.

Table 3—Results of Different Targeting Methods on Error Rate Based on Consumption

Sample:	Full population (1)	By income status		By detailed income status				Per capita consumption of beneficiaries (8)
		Inclusion error (2)	Exclusion error (3)	Rich (4)	Middle income (5)	Near poor (6)	Very poor (7)	
Community treatment	0.031*	0.046**	0.022	0.028	0.067**	0.49	−0.013	9.933
	(0.017)	(0.018)	(0.028)	(0.021)	(0.027)	(0.038)	(0.039)	(18.742)
Hybrid treatment	0.029*	0.037**	0.009	0.020	0.052**	0.031	−0.008	−1.155
	(0.016)	(0.017)	(0.027)	(0.020)	(0.025)	(0.037)	(0.037)	(19.302)
Observations	5,753	3,725	2,028	1,843	1,882	1,074	954	1,719
Mean in PMT treatment	0.30	0.18	0.52	0.13	0.23	0.55	0.48	366

Notes: All regressions include stratum fixed effects. Robust standard errors in parentheses, clustered at the village level. All coefficients are interpretable relative to the PMT treatment, which is the omitted category. The mean of the dependent variable in the PMT treatment is shown in the bottom row. All specifications include stratum fixed effects.

   ***Significant at the 1 percent level.
   **Significant at the 5 percent level.
    *Significant at the 10 percent level.

erty lines based on the PPP\$2 per-day consumption threshold, and then classify a household as incorrectly targeted if its per capita consumption level is below the poverty line and it was not chosen as a beneficiary, or if it was above the poverty line and it was identified as a recipient ($Error_{ivk}$). We then examine which method minimized the error rate by estimating the following equation using OLS:

$$(1) \quad ERROR_{ivk} = \alpha + \beta_1 COMMUNITY_{ivk} + \beta_2 HYBRID_{ivk} + \gamma_k + \varepsilon_{ivk},$$

where i represents a household, v represents a subvillage, k represents a stratum, and γ_k are stratum fixed effects.[23] Note that the PMT treatment is the omitted category, so β_1 and β_2 are interpretable as the impact of the community and the hybrid treatments relative to the PMT treatment. Since the targeting methods were assigned at the subvillage level, the standard errors are clustered to allow for arbitrary correlation within a subvillage.

The results, shown in Table 3, indicate that the PMT method outperforms both the community and hybrid treatment in terms of the consumption-based error rate. Under the PMT, 30 percent of the households are incorrectly targeted (column 1).[24] Both the community and hybrid methods increase the error rate by about three percentage points—or about ten percent—relative to the PMT method (significant at the ten percent level).[25]

---

[23] For simplicity of interpretation, we use OLS/linear probability models for all dependent variables in Table 3. Using a probit model for the binary dependent variables produces results of the same sign and significance level.

[24] Fluctuations in consumption between the date of the baseline survey and that of targeting could lead to overinflated error rates. To minimize this, we ensured that the targeting quickly followed the baseline survey: the average time lapse was 44 days. We also ensured that the time between the baseline survey and the targeting was orthogonal to the treatment. Online Appendix Table 3 shows that the time between survey and targeting date has no effect on the error rates, and that the interaction of time elapsed with the treatment dummies is never significant.

[25] The community treatment does not provide any indication of the absolute level of poverty. Thus, we chose the fraction of households in each subvillage that would become beneficiaries through geographic targeting. For consistency, we use geographic targeting across all three treatments. By imposing this constraint on the PMT, however, we do not take full advantage of the fact that it provides absolute measures of poverty. Taking advantage of
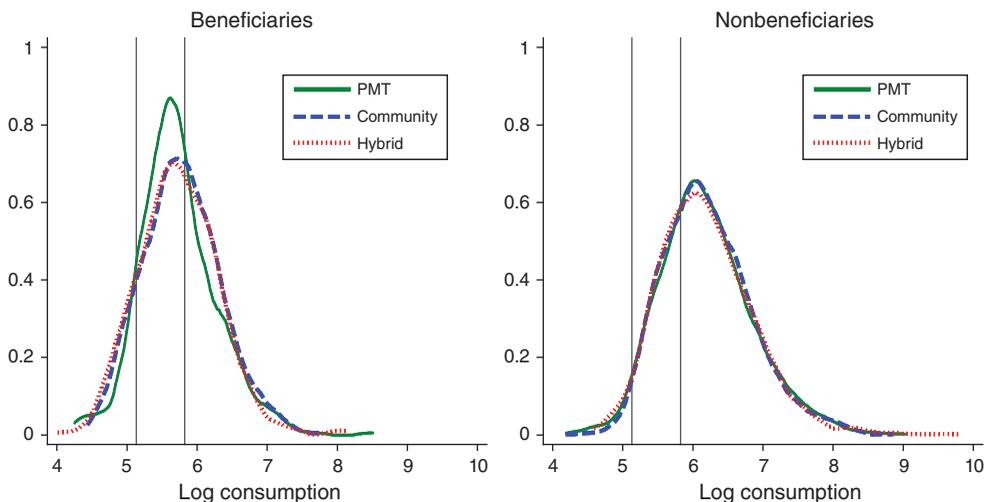
FIGURE 1. PDF OF LOG PER CAPITA CONSUMPTION OF BENEFICIARIES AND NONBENEFICIARIES, BY TREATMENT STATUS

Notes: The left panel shows the Probability Density Function (PDF) of log per capita consumption for those house-holds chosen to receive the transfer, separately by each treatment. The right panel shows the PDF of log per cap-ita consumption for those households not chosen to receive the transfer, separately by treatment. The vertical lines show PPP$1 and PPP$2 per day poverty lines (see footnote 19 for more information on the calculation of these poverty lines).

Adding a rich household to the list may have different welfare implications than adding a household that is just above the poverty line. To examine this, Figure 1 graphs the log per capita consumption distribution of the beneficiaries (left panel) and nonbeneficiaries (right panel) for each targeting treatment. The vertical lines in the graphs indicate PPP$1 and PPP$2 per-day poverty lines. Overall, the graphs confirm that all methods select relatively poorer households: for all methods, the mode per capita consumption for beneficiaries is below PPP$2 per day, whereas it is above PPP$2 per day for nonbeneficiaries.

Examining the impact of the treatments, the left panel shows that the consumption distribution of beneficiaries derived from the PMT is centered to the left of the distribution under the community and hybrid methods. Thus, on average, the PMT identifies poorer individuals. The community methods, however, select a greater percentage of beneficiaries whose log daily per capita consumption is less than PPP$1 (the leftmost part of the distribution). Thus, the figures suggest that despite doing worse on average, the community methods may capture more of the *very* poor. Moreover, the figures suggest that all three methods contain similar propor-tions of richer individuals (with log income greater than about 6.5). The difference in the error rate across the three treatments is driven by differences in the near poor (PPP$1 to PPP$2) and the middle-income group (those above the PPP$2 poverty line, but with log income less than 6.5).

We more formally examine the findings from Figure 1 in the remaining columns of Table 3. In columns 2 and 3, we examine the error rates separately for the poor (exclu-

_____

this information, the PMT would perform 6 percentage points (or 20 percent) better than the community methods in selecting the poor. This analysis is available upon request.

sion error) and the nonpoor (inclusion error). In columns 4 and 5, we disaggregate the nonpoor into rich and middle, and in columns 6 and 7, we disaggregate the poor by splitting them into near poor and very poor. The results confirm that much of the difference in the error rate between the community methods and the PMT occurs near the cutoff for inclusion. Specifically, the community and hybrid methods are, respectively, 6.7 and 5.2 percentage points more likely to misclassify the middle nonpoor (column 5, both statistically significant at five percent). They are also more likely to misclassify the near poor by 4.9 and 3.1 percentage points, respectively, although these results are not individually statistically significant. In contrast, we observe much less difference between the methods for the rich and the very poor, and in fact the point estimate suggests that the community method may actually do better among the very poor.

In column 8, we examine the average per capita consumption of beneficiaries across the three groups. As expected, given that the community treatment selects more of the very poor and also selects more individuals who are just above the PPP\$2 poverty line, the average per capita consumption of beneficiaries is not substantially different between the various treatments. This suggests that even though the community treatments are more likely to mistarget the poor as defined by the PPP\$2 cutoff, the welfare implications of the three methods appear similar based on the consumption metric.[26]

Given that the levels of information and capture may be different across localities, we examine the heterogeneity in the relative effectiveness of the different treatments across three dimensions, all of which we specified ex ante when designing the intervention. First, we hypothesized that the community methods may do worse in urban areas, where individuals may not know their neighbors as well. Our sample was stratified along this dimension to ensure that we had a large enough sample size to test this hypothesis. Second, the level of inequality in the villages could result in important differences between the two techniques. On one hand, community-based targeting may work better in areas with large inequality, since it implies that the rich and the poor are more sharply differentiated. On the other hand, elite capture of community-based techniques may be more severe in areas with high inequality if rich elites are powerful enough to exclude the poor from the community decision-making process. Third, we hypothesized that in the areas where many people are related to one other by blood or marriage, they have more information about their neighbors, so the community method should work better.

We present the results of the analysis where we interact the various treatment variables with these three dimensions of heterogeneity in Table 4.[27] We find that, in general, the error rate was *lower* in the community treatment (relative to the PMT) in urban areas, in areas with high inequality, and in areas where many households are related. These effects are not significant at conventional levels, however. In addition, we also test whether the treatments differed in Java and the other provinces, as

---

[26] To maximize social welfare, the targeting method should select households with the highest average marginal utility. If utility is quadratic in per capita consumption, marginal utility is exactly equal to per capita consumption, so the regression in column 8 shows that there is no difference in average marginal utility across the three treatments based on this metric. In results not reported in the table, we have also confirmed that the average marginal utility of beneficiaries is the same across treatments using alternate specifications for the utility function as well, including CRRA utility with $\rho = 1$ (log), 2, 3, 4, and 5.

[27] Note that we define inequality as the range between the twentieth and the eightieth percentile per capita consumption levels.

Table 4—Results of Different Targeting Methods on Error-Heterogeneity

Sample:	Full population (1)	By income status Exclusion error (2)	Inclusion error (3)	By detailed income status Rich (4)	Middle income (5)	Near poor (6)	Very poor (7)	Per capita consumption of beneficiaries (8)
Community treatment	0.069**	0.005	0.145**	−0.052	0.068	0.218***	0.042	−44.804
	(0.035)	(0.039)	(0.058)	(0.050)	(0.053)	(0.079)	(0.083)	(36.192)
Hybrid treatment	0.087**	0.017	0.130**	0.041	−0.009	0.200**	0.092	−22.408
	(0.037)	(0.042)	(0.060)	(0.054)	(0.054)	(0.078)	(0.087)	(40.155)
Urban village	−0.010	−0.098***	0.128***	−0.088**	−0.113**	0.231***	0.035	−20.668
	(0.030)	(0.035)	(0.046)	(0.043)	(0.047)	(0.063)	(0.062)	(36.623)
Inequality	−0.004	−0.026	0.057	−0.029	−0.015	0.043	0.091*	−3.963
	(0.026)	(0.027)	(0.040)	(0.033)	(0.039)	(0.055)	(0.053)	(33.960)
General connectedness	0.043*	−0.010	0.000	0.009	−0.046	0.036	−0.043	−20.957
	(0.026)	(0.031)	(0.039)	(0.039)	(0.039)	(0.053)	(0.056)	(29.451)
Urban village × community treatment	−0.049	0.026	−0.110*	0.056	−0.029	−0.184**	−0.041	25.062
	(0.036)	(0.039)	(0.060)	(0.045)	(0.058)	(0.078)	(0.083)	(41.936)
Urban village × hybrid treatment	−0.051	0.014	−0.069	−0.001	0.020	−0.149*	−0.015	11.794
	(0.036)	(0.039)	(0.059)	(0.045)	(0.055)	(0.079)	(0.082)	(47.509)
Inequality × community treatment	−0.016	0.022	−0.134**	0.027	0.052	−0.128*	−0.117	38.764
	(0.035)	(0.039)	(0.055)	(0.047)	(0.057)	(0.075)	(0.076)	(39.728)
Inequality × hybrid treatment	−0.021	−0.002	−0.069	−0.036	0.061	0.006	−0.157**	13.022
	(0.034)	(0.037)	(0.054)	(0.045)	(0.054)	(0.073)	(0.077)	(41.603)
General connectedness × community treatment	−0.018	0.022	−0.007	0.069	−0.025	−0.038	0.040	44.665
	(0.035)	(0.041)	(0.058)	(0.050)	(0.056)	(0.078)	(0.082)	(39.438)
General connectedness × hybrid treatment	−0.050	0.018	−0.081	−0.007	0.034	−0.166**	−0.010	17.931
	(0.035)	(0.040)	(0.058)	(0.050)	(0.052)	(0.080)	(0.078)	(40.290)
Observations	5,753	3,725	2,028	1,843	1,882	1,074	954	1,719
Mean in PMT treatment	0.30	0.18	0.52	0.13	0.23	0.55	0.48	366

Note: See notes to Table 3.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

previous studies (e.g., Ravallion and Dearden 1988) have shown that Java tends to be more egalitarian (note that our sample was also stratified along this dimension). The results of this analysis, presented in online Appendix Table 4, show no substantive differences between Java and the other provinces.

In sum, we do not observe significant differences between the methods based on the four levels of heterogeneity that we considered. This speaks to the external validity of the study, suggesting that the findings may be easier to generalize to other settings.

B. *Effects of Targeting Policy on Poverty Rate and Gap*

We observe that the community treatment has a three percentage point–larger error rate based on consumption than the PMT. Given that the differences are largely driven by those near the thresholds, an important question is whether this is large enough to affect real outcomes, such as the head count poverty rate (the percentage of people who fall below the poverty line) and the poverty gap (the mean distance below the poverty line as a proportion of the line, counting the nonpoor as having zero gap). Moreover, given that the community method better targets the very poor, it is possible that the community methods may perform better at reducing the squared

TABLE 5—SIMULATED POVERTY IMPACTS FOR DIFFERENT TRANSFER SIZES AND POVERTY LINES

Transfer size (Rp. 1,000s)		Poverty line = poor			Poverty line = very poor		
		PMT	Community	Hybrid	PMT	Community	Hybrid
No transfer	Head count	33.86	33.86	33.86	15.64	15.64	15.64
	Pov. gap	9.45	9.45	9.45	3.55	3.55	3.55
	Sq. pov gap	3.73	3.73	3.73	1.21	1.21	1.21
50	Head count	32.81	32.81	33.30	15.04	14.52	14.55
	Pov. gap	8.88	8.89	8.87	3.17	3.15	3.16
	Sq. pov gap	3.42	3.41	3.39	1.04	1.02	1.02
100	Head count	31.45	32.03	32.01	14.02	13.68	13.77
	Pov. gap	8.37	8.40	8.35	2.84	2.82	2.83
	Sq. pov gap	3.14	3.13	3.11	0.91	0.88	0.88
200	Head count	29.60	30.03	30.10	12.33	12.13	12.28
	Pov. gap	7.47	7.54	7.48	2.37	2.33	2.32
	Sq. pov gap	2.71	2.69	2.67	0.74	0.68	0.68
500	Head count	24.22	25.15	25.38	8.85	9.17	9.13
	Pov. gap	5.76	5.87	5.84	1.69	1.59	1.61
	Sq. pov gap	2.02	1.98	1.97	0.53	0.45	0.46

Notes: Transfers are in thousands of rupiah (Rp. 1,000 = US $0.10) and assume a uniform transfer for all recipient households. Note that the transfer is per household, whereas poverty is defined as per capita, so a given transfer has less of an impact on larger households. Head count and poverty gap are in percent; squared poverty gap is multiplied by 100.

poverty gap (which places greater weight on reducing the poverty of the very poor), even if it performs worse in reducing the poverty head count ratio.

We follow the methods used in Ravallion (2009) and simulate the effects of the different targeting methods on the head count poverty rate, the poverty gap, and squared poverty gap. We provide the results of the simulation for selected transfer sizes (no transfer, Rp. 50,000, Rp. 100,000, Rp. 200,000, and Rp. 500,000 per month) in Table 5, and graph out the full results by transfer size in Figure 2. In both Table 5 and Figure 2, we focus on the poor and very poor poverty lines. Note that despite the randomization, there are statistically insignificant differences between the poverty rates in the different treatments as a result of sampling; for the simulations, we assume for all treatments the distribution of consumption from the PMT villages, so that we have exactly the same income distribution across treatments.[28]

The differences in targeting accuracy across the three methods do not result in large differences in the measures of poverty under consideration. For example, Indonesia's BLT program provided transfers of Rp. 100,000 per month for a year. If a similar transfer was provided to those below the poor poverty line, the PMT would reduce the head count ratio from 33.86 to 31.45 (column 1), while the community and hybrid methods would result in ratios of 32.03 and 32.01, respectively (columns 2 and 3).[29]

[28] We first weight households so that the weighted distribution of households has exactly the same number of beneficiaries in each treatment group. Second, within each treatment group, we compute each household's weighted percentile rank in the per capita consumption distribution, and assign that household the per capita consumption of the corresponding household from the PMT treatment. These two very minor adjustments correct for small-sample differences in the underlying consumption distribution between the three treatment groups and ensure that the only differences in Table 5 are due to the differences in targeting outcomes.

[29] To test for whether these differences are statistically significant, we assign each household a variable equal to their contribution to the poverty metric (i.e., if the poverty metric is the head count ratio, we assign that household a 1 if it is below the poverty line after the transfer and 0 otherwise; likewise for poverty gap and squared poverty gap). We

Panel A. Head count ratio
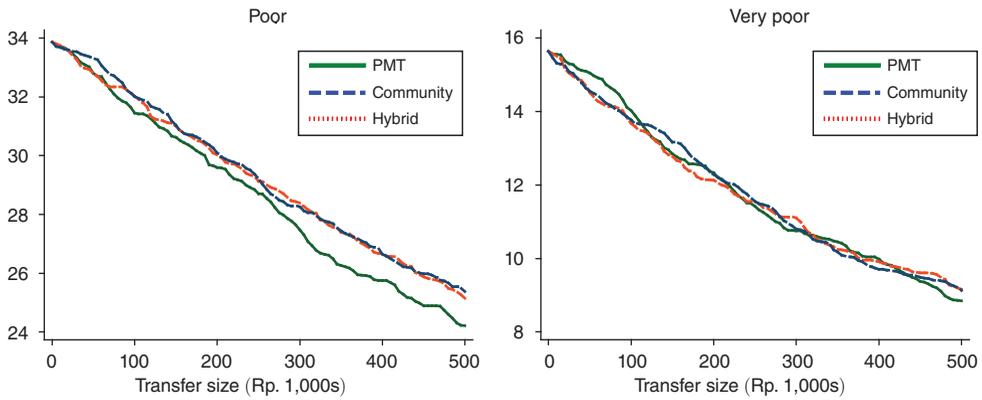
Panel B. Poverty gap
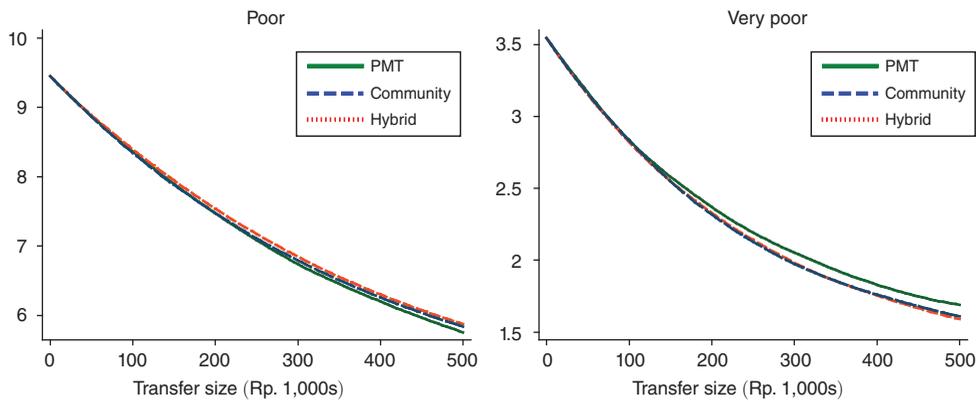


Panel C. Squared poverty gap
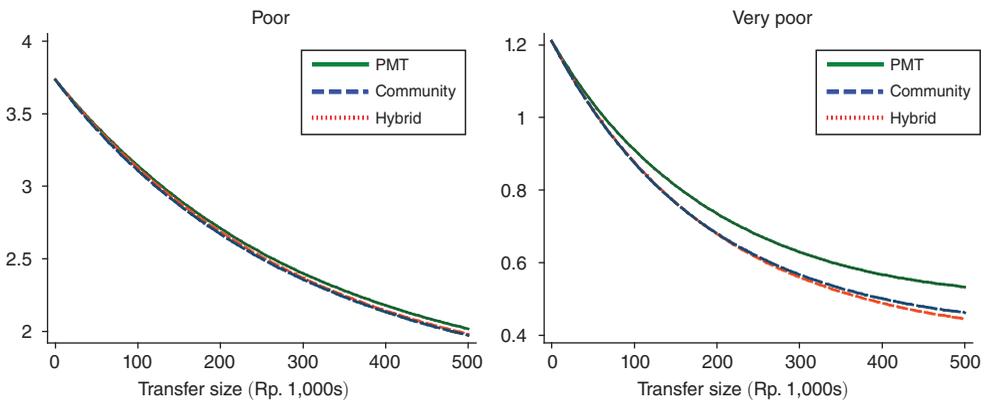
FIGURE 2. SIMULATED POVERTY IMPACTS FOR DIFFERENT TRANSFER SIZES AND POVERTY LINES

Notes: These figures show the results of simulations of the poverty impact of targeting schemes of alternative size using each of the three targeting methods. In panel A, the y-axis shows the head count ratio; in panel B, it shows the poverty gap; and in panel C, it shows the squared poverty gap, using either the poor (left column) or very poor (right column) poverty lines. The x-axis in each graph shows the size of the transfer.

---

then run a regression of the poverty metric (at the household level) on a dummy for the treatment group, and cluster standard errors at the village level. None of the differences in the table are statistically significant using this metric.

The poverty gap for the three methods is also not significantly different (7.47 for PMT, 7.54 for the community method, and 7.48 for the hybrid method). Given that the community method works better at identifying the very poor, however, the community treatments actually do better at reducing the poverty head count (from 15.64 to 13.68 for the community and to 13.77 for the hybrid) than the PMT (14.02) at the very poor poverty line, but this difference is not statistically significant. The squared poverty gap, with its emphasis on the very poorest, gives the best chance for the community method to dominate the PMT. With the very poor poverty line, for a transfer of Rp 100,000 per month, the squared poverty gaps are still quite similar for the PMT (0.91) and the community method (0.88). Doubling the transfer results in the community method doing substantially better—0.68 for the community versus 0.74 for the PMT, though the difference is still not statistically significant.

Note that these baseline simulations do not include differences in targeting costs. To account for the costs of targeting, we assume that the targeting is done annually. We therefore divide the targeting costs by the number of households per village who are beneficiaries to provide a per-beneficiary cost, and divide by 12 to obtain a monthly cost. We then reduce the transfer by this amount. Since the targeting costs are small when expressed monthly in this way (the costs are Rp. 7,000 per beneficiary per month for PMT, compared with Rp. 3,100 for community, and Rp. 8,000 for hybrid), they do not qualitatively affect the results above. Online Appendix Table 5 reproduces Table 5 after subtracting out these targeting costs.

C. Satisfaction

In Table 6, we study the impacts of the treatments on the communities' satisfaction levels and the legitimacy of the targeting. Panel A presents data from the endline household survey. Panel B presents data from the follow-up survey of subvillage heads. Panel C presents the results from the anonymous comment box, the community's complaints to the village head, and the facilitator comments on the ease of distributing the transfer payments.[30]

Individuals are much more satisfied with the community treatment than with the PMT or hybrid treatments (panel A). For example, in the community treatment, respondents wish to make fewer changes to the beneficiary list; they would prefer to add about ⅓ fewer households to the list of beneficiaries (column 4) and subtract about ½ as many households (column 5) than in the PMT or the hybrid treatments. Individuals in the community treatment are more likely to report that the method was appropriate (column 1) and are also more likely to state that they are satisfied with the program (column 2). A joint test of these dependent variables indicates that the community treatment differences are jointly statistically significant (p-value < 0.001).

Subvillage heads are also much more satisfied (panel B). The subvillage head was 38 percentage points more likely to say that the targeting method was appropriate when community-based targeting was used and 17 percentage points less

---

[30]For simplicity of interpretation, we use OLS/linear probability models for all dependent variables in Table 6. Using ordered probit for categorical response variables and probit for binary dependent variables produces the same signs of the results, and the same levels of statistical significance.

TABLE 6—SATISFACTION

Panel A. Household endline survey

	Is the method applied to determine the targeted households appropriate? (1 = worst, 4 = best) (1)	Are you satisfied with the targeting activities in this subvillage in general? (1 = worst, 4 = best) (2)	Are there any poor HH that should be added to the list? (0 = no, 1 = yes) (3)	Number of HH that should be added to list (4)	Number of HH that should be subtracted from list (5)	*p*-value from joint test (6)
Community treatment	0.161*** (0.056)	0.245*** (0.049)	−0.189*** (0.040)	−0.578*** (0.158)	−0.554*** (0.112)	< 0.001
Hybrid treatment	0.018 (0.055)	0.063 (0.049)	0.020 (0.042)	0.078 (0.188)	−0.171 (0.129)	0.762
Observations	1,089	1,214	1,435	1,435	1,435	
Mean in PMT treatment	3.243	3.042	0.568	1.458	0.968	

Panel B. Subvillage head endline survey

	Is the method applied to determine the targeted households appropriate? (0 = no, 1 = yes)	In your opinion, are villagers satisfied with the targeting activities in this subvillage in general? (1 = worst, 4 = best)	Are there any poor HH that should be added to the list? (0 = no, 1 = yes)	Are there any nonpoor HH that should be subtracted from the list? (0 = no, 1 = yes)	
Community treatment	0.378*** (0.038)	0.943*** (0.072)	−0.169*** (0.045)	−0.010 (0.020)	< 0.001
Hybrid treatment	0.190*** (0.038)	0.528*** (0.071)	−0.065 (0.043)	−0.019 (0.019)	< 0.001
Observations	636	629	640	640	
Mean in PMT treatment	0.565	2.456	0.732	0.057	

Panel C. Comment forms and fund disbursement results

	Number of comments in the comment box	Number of complaints in the comment box	Number of complaints received by subvillage head	Did facilitator encounter any difficulty in distributing the funds? (0 = no, 1 = yes)	Fund distributed in a meeting (0 = no, 1 = yes)	
Community treatment	−0.944 (0.822)	−1.085*** (0.286)	−2.684*** (0.530)	−0.062*** (0.023)	0.082** (0.038)	0.0014 0.177
Hybrid treatment	−0.364 (0.821)	−0.554** (0.285)	−2.010*** (0.529)	−0.045* (0.026)	0.051 (0.038)	
Observations	640	640	640	621	614	
Mean in PMT treatment	11.392	1.694	4.34	0.135	0.579	

Notes: All estimation is by OLS with stratum fixed effects. Using ordered probit for multiple response and probit models for binary dependent variables produces the same signs and statistical significance as the results shown. These results are available from the authors upon request.

likely to name any households that should be added to the list. The higher levels of satisfaction were also manifested in fewer complaints (panel C). There were on average 1.09 fewer complaints in the comment box in the community subvillages relative to the PMT subvillages, and 0.55 fewer complaints in the hybrid subvillages relative to the PMT (column 2). The subvillage head also reported receiving

2.68 and 2.01 fewer complaints in the community and hybrid treatment, respectively (column 3).

The higher satisfaction levels in the community treatment led to a smoother disbursal process. First, the facilitators who distributed the cash payment were four to six percentage points less likely to experience difficulties while doing so in subvillages assigned to the community or hybrid method (panel C, column 4). Second, the subvillage heads could choose for the cash disbursals to happen in an open community meeting or, if the head felt that they would encounter problems in the village, the facilitator could distribute the transfer door to door. Facilitators were eight percentage points more likely to distribute the cash in an open meeting in the subvillages assigned to the community treatment (panel C, column 5).[31]

D. *Understanding the Differences between PMT and Community Targeting*

The findings present an interesting puzzle. The results on the error rates suggest that the community-based methods actually do somewhat worse at identifying the poor, although this does not impact the poverty rate significantly. The community method results in greater satisfaction levels, however. The following sections explore alternative explanations of why the PMT and the community methods differ: elite capture, community effort problems, heterogeneity in preferences within the villages, and differences in information.

IV. Elite Capture

In Table 7, we test for elite capture by examining whether elite connected households are more likely to be beneficiaries when the elites potentially have more control of the process (i.e., in the elite-only meetings). We first verify that the elite-only meetings had an impact on attendance, and then test for whether the elite-only treatment affected the error rate. We reestimate equation (1) (with both attendance measures and the error rate as outcome variables), but now include a dummy for the *ELITE* subtreatment. As expected, elite meetings have lower participation: we find that 48 percent of households attended the meetings in the

---

[31] Do these differences in satisfaction represent changes from the act of directly participating in the process (as in Olken 2010), from knowing that some local process was followed, or from changes in the final list of beneficiaries? We find no differences in our measures of satisfaction between the whole community treatments (when 48 percent of households attended the meeting) and elite community treatments (when only 17.6 percent of households attended the meeting), suggesting that either differences in the list or knowing that a local process was followed drives the differences in satisfaction. It is hard to differentiate between these remaining two hypotheses. To test for whether differences in satisfaction were driving the results, we computed an approximate PMT score for each individual in the baseline and then computed the rank correlation between this score and the targeting rank list that resulted from the experiment. We created a dummy variable that indicates a high correlation between these two measures, and interacted this variable with the community and hybrid treatments. There is no discernible difference across the satisfaction measures, implying that the higher satisfaction that was observed in the community treatment was not affected by the degree to which the community's list would match the PMT. This suggests that knowing that a local process was followed seems to drive the satisfaction levels. These results are available upon request. On the other hand, as discussed above, when shown the resulting beneficiary list, community members made fewer additions and subtractions to the list, suggesting that they actually have fewer disagreements with the resulting names in the community treatment, so the difference in the list itself may also be important.

TABLE 7—ELITE TREATMENTS

	Attendance (survey data)	Full sample error rate	Full sample error rate		On beneficiary list	
	(1)	(2)	(3)	(4)	(5)	(6)
Community treatment	0.367***	0.029	0.033	0.048*	0.042*	0.054*
	(0.038)	(0.018)	(0.023)	(0.025)	(0.025)	(0.028)
Hybrid treatment	0.370***	0.027	0.024	0.008	0.025	0.012
	(0.037)	(0.018)	(0.022)	(0.024)	(0.022)	(0.023)
Elite subtreatment	−0.301***	0.004	0.016	−0.013	−0.015	−0.039
	(0.034)	(0.016)	(0.020)	(0.029)	(0.021)	(0.032)
Elite × hybrid				0.062		0.051
				(0.041)		(0.043)
Elite connectedness			−0.025	−0.025	−0.063***	−0.063***
			(0.021)	(0.021)	(0.021)	(0.021)
Elite connectedness × community treatment			−0.015	−0.013	−0.067**	−0.078**
			(0.035)	(0.038)	(0.033)	(0.036)
Elite connectedness × hybrid treatment			0.010	0.010	−0.013	−0.001
			(0.033)	(0.035)	(0.033)	(0.035)
Elite connectedness × elite treatment			−0.029	−0.034	0.041	0.064
			(0.031)	(0.047)	(0.030)	(0.042)
Elite connectedness × elite treatment × hybrid				0.003		−0.047
				(0.063)		(0.060)
Observations	287	5,753	5,753	5,753	5,756	5,756
Mean in PMT treatment	0.11	0.30	0.30	0.30	0.28	0.28

Notes: In column 1, an observation is a village and the dependent variable is the share of households surveyed in the endline survey where at least one household member attended a targeting meeting. The PMT mean in column 3 is not zero, because the question was worded generically to be about any targeting meeting, not just meetings associated with our project. The dependent variable in columns 2–4 is the dummy for error in targeting based on consumption, as in column 1 of Table 3. Dependent variable in columns 5 and 6 is a dummy for being a beneficiary of the program. All specifications in columns 3–6 include dummies for the community, hybrid, and elite treatment main effects, as well as stratum fixed effects; columns 4 and 6 also include a dummy for elite × hybrid. Robust standard errors are in parentheses, and standard errors are adjusted for clustering at the village level in columns 2–6. All specifications include stratum fixed effects.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

whole community treatment, compared to 18 percent in the elite subtreatment (column 1).[32] The error rate, however, was not significantly different across the two treatments (column 2).

While the evidence presented in column 2 is consistent with no elite capture, it is also consistent with the elite dominating the whole community meeting, leading to the result that both types of meetings reflect their preferences.[33] To test this, we examine whether the elites and their relatives were more likely to be selected in

---

[32] The data on attendance come from questions about generic targeting meetings during the endline survey. It is possible to report having attended a meeting (such as a meeting during the socialization of the program or a meeting about another targeted related activity) even in villages where our project held no ranking meeting.

[33] This second story seems unlikely: the facilitators report that a few individuals dominated the conversation in only 15 percent of the meetings, and that otherwise the meetings were full community affairs.

both the whole community and elite meetings relative to the PMT. Specifically, we estimate the following equation:

$$(2) \quad ERROR_{ivk} = \alpha + \beta_1 COMMUNITY_{ivk} + \beta_2 HYBRID_{ivk} + \beta_3 ELITE_{ivk}$$

$$+ \beta_4 CONN_{ivk} + \beta_5 (COMMUNITY_{ivk} \times CONN_{ivk})$$

$$+ \beta_6 (HYBRID_{ivk} \times CONN_{ivk}) + \beta_7 (ELITE_{ivk} \times CONN_{ivk})$$

$$+ \gamma_k + \varepsilon_{ivk},$$

where $CONN_{ivk}$ is an indicator that equals one if the household is related to any of the subvillage leaders/elites, or is one of the leaders themselves.[34] Columns 3 and 4 examine the error rate as the dependent variable, and columns 5 and 6 examine whether a household received the transfer as the dependent variable. We find little evidence of elite capture. In fact, the point estimates suggest the opposite: the elite-connected households are less likely to be mistargeted in the community and elite treatments, although the effect is not significant at conventional levels. In fact, we find that the elites are actually penalized in the community treatments: elites and their relatives are about 6.7 to 7.8 percent less likely to be on the beneficiary list in the community treatments relative to PMT meetings.[35] The point estimates suggest that the elite treatment undoes this penalty to some degree, but on net in elite versions of the community meetings, elites are still 2.6 percentage points less likely to receive transfers than in the PMT treatment (though the combined effect is not statistically significant from 0). These findings suggest that elite capture is not the reason that the mistargeting is worse under the community method.

V. Problems with Community Effort

The community-based ranking process requires human effort: ranking 75 households would require making at least 363 pairwise comparisons. Thus, the worse targeting in the community methods could result simply from fatigue as the ranking progresses. To investigate this, we randomized the order in which households were ranked.

Figure 3 graphs the relationship between the error rate and the randomized rank order from a nonparametric Fan regression, with cluster-bootstrapped 95 percent

[34] Specifically, we defined an "elite connected" household as any household where (i) we interviewed the household and found that a household member held a formal leadership position in the village, such as village or subvillage head; (ii) at least two of the respondents we interviewed identified the household as holding either a formal or informal (*tokoh*) leadership role in the village; or (iii) a household connected by blood or marriage to any household identified in (i) or (ii).

[35] It is possible that the elite connected households are more likely to be connected to other households in the subvillage in general. In this case, the penalty in columns 3 and 4 may not be due to the fact that they are elite, but instead to the fact that the community believes that they will be "taken care of" by their relatives. In online Appendix Table 5, we rerun the specifications in Table 6 (columns 3–6), now including both the main effects and interacted effects of the households' general connectedness within the village (specifically, a dummy variable for whether the household is related by blood or marriage to any other household in the village) as well as elite connectedness. The elite results stay robust (both in magnitude and significance) when controlling for general connectedness.
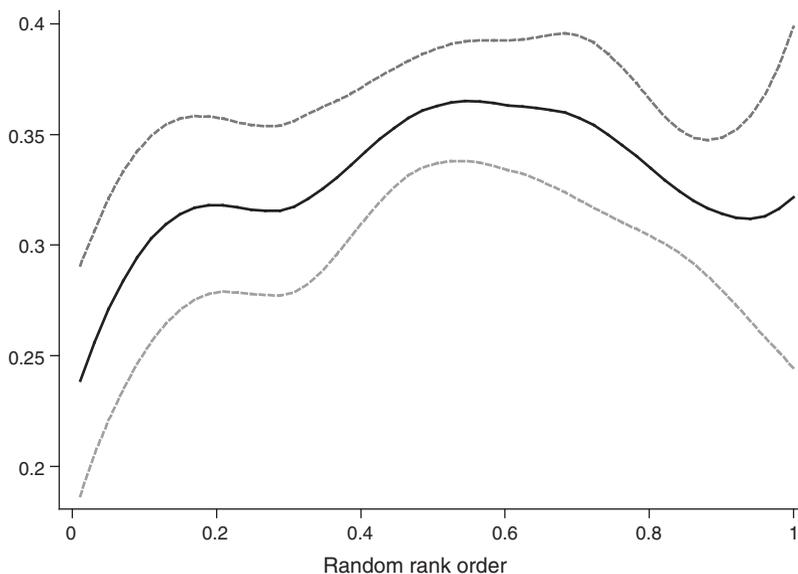
FIGURE 3. EFFECT OF ORDER IN RANKING MEETING ON MISTARGET RATE

Notes: This figure graphs the relationship between mistargeting and the randomized rank order from a nonparametric Fan regression. The dashed lines represent cluster-bootstrapped 95th percentile confidence intervals.

confidence intervals shown as dashed lines. The error rate is lowest for the first few households ranked, but then rises sharply by the twentieth percentile of households. The magnitude is substantial—the point estimates imply that error rates are between five to ten percentage points lower for the first household than for households ranked in the latter half of the meeting.

Table 8 reports results from investigating these issues in a regression framework. Column 1 reports the results from estimating the relationship between the error rate and the randomized rank order, which varies from 0 (household was ranked first) to 1 (the household was ranked last). The point estimate is positive, indicating a higher error rate for households ranked later, but it is not statistically significant. In column 2, we interact the order with the hybrid treatment. The results show that in the community treatment, there is substantially more error at the end of the process: the first household ranked is 5.6 percentage points less likely to be incorrectly targeted than the last household ranked; these results are marginally significant with a *p*-value of 0.11. On net, the community treatment actually does slightly better than the PMT in the beginning, but substantially worse towards the end. This effect is completely undone in the hybrid, where the random rank order and the error rate appear unrelated. Columns 3 and 4 examine how the rank order affects whether a household receives the transfer. On average, households ranked at the end of the meeting are 4.9 percentage points more likely to be on the beneficiary list than those ranked at the start (significant at the ten percent level). The additional error from being late in the list thus comes largely from richer households ranked toward the end of the process being more likely to be on the list.

TABLE 8—EFFORT

	Mistarget dummy		On beneficiary list dummy	
	(1)	(2)	(3)	(4)
Household order in ranking (percentile)	0.029 (0.026)	0.056 (0.037)	0.049* (0.026)	0.048* (0.029)
Household order in ranking × hybrid		−0.053 (0.052)		0.001 (0.028)
Observations	3,784	3,784	3,785	3,785

Notes: All specifications are limited to community and hybrid villages. Columns 1 and 2 include a hybrid dummy and stratum fixed effects; columns 3 and 4 include stratum fixed effects since the total number of beneficiaries is constant in all treatments. The dependent variable in columns 1 and 2 is the mistarget dummy for the full sample, as in column 1 of Table 4. The dependent variable in columns 3 and 4 is a dummy for being chosen as a recipient.

　***Significant at the 1 percent level.
　**Significant at the 5 percent level.
　　*Significant at the 10 percent level.

VI. Does the Community Have a Different Maximand?

A third potential reason why the community produced a different outcome than the PMT is that the community is doing its best to identify the poor, but has a different concept of poverty. We thus explore whether the community's views on poverty differ from that of per capita consumption.

A. *Alternative Welfare Metrics*

We begin by examining how the targeting outcomes compare not just against the government's metric of welfare (captured by r_g, the ranking based on per capita consumption), but also against alternative welfare metrics. In our baseline survey, we asked eight randomly chosen members of the community to confidentially rank each other from poorest to richest. We average the ranks to construct each household's wealth rank according to the other community members, denoted r_c. To capture welfare as measured from an elite perspective, denoted r_e, we examine how the subvillage head ranked these eight other households. To measure how people assess their own poverty, denoted r_s, we asked all respondents to rate their own poverty level on a scale of 1 to 6. We computed the percentile rank of each measure to put them on the same scale.[36]

To assess the poverty targeting results against these alternative welfare metrics, we compute the rank correlation between the targeting rank list derived from the experiment and each of four welfare metrics. We then examine the effectiveness

---

[36] Online Appendix Table 7 presents the matrix of rank correlations between these alternative welfare metrics. The correlation matrix shows that while all of the welfare metrics are positively correlated, they clearly capture different things. Of particular note is the self-assessments: while the rank correlation of self-assessments (r_s) with consumption (r_g) is only 0.26, that with community survey ranks (r_c) is 0.45. Thus, the survey ranks appear to capture how individuals feel about themselves better than per capita consumption.

TABLE 9—ASSESSING TARGETING TREATMENTS USING ALTERNATIVE WELFARE METRICS

	Consumption (r_g) (1)	Community survey ranks (r_c) (2)	Subvillage head survey ranks(r_e) (3)	Self-assessment (r_s) (4)
Community treatment	−0.065** (0.033)	0.246*** (0.029)	0.248*** (0.038)	0.102*** (0.033)
Hybrid treatment	−0.067** (0.033)	0.143*** (0.029)	0.128*** (0.038)	0.075** (0.033)
Observations	640	640	640	637
Mean in PMT treatment	0.451	0.506	0.456	0.343

Notes: The dependent variable is the rank correlation between the treatment outcome (i.e., the rank ordering of households generated by the PMT, community, or hybrid treatment) and the welfare metric shown in the column, where each observation is a village. Robust standard errors are shown in parentheses.

   ***Significant at the 1 percent level.
   **Significant at the 5 percent level.
    *Significant at the 10 percent level.

of the various targeting treatments against these different measures of well-being by estimating

$$(3) \quad RANKCORR_{vkR} = \alpha + \beta_1 COMMUNITY_{vk} + \beta_2 HYBRID_{vk} + \gamma_k + \varepsilon_{vkR},$$

where $RANKCORR_{vkR}$ is the rank correlation between the targeting rank list and the well-being measure R in subvillage v. Stratum fixed effects (γ_k) are included. The results are reported in Table 9. As the data is aggregated to the village level, each regression has 640 observations.

The results provide striking evidence that per capita consumption as we measure it does not fully capture what the community calls welfare. Column 1 confirms the results that are shown in Table 3: both the community and hybrid treatment result in lower rank correlations with per capita consumption than the PMT. Specifically, they are 6.5–6.7 percentage points, or about 14 percent, lower than the rank correlations obtained with PMT. They move away from consumption in a very clear direction, however—the community treatment increases the rank correlation with r_c by 24.6 percentage points, or 49 percent above the PMT level. The hybrid also increases the correlation with r_c but the magnitude is about half that of the community treatment. Thus, the verification in the hybrid appears to move the final outcome away from the community's perception of well-being. These differences are statistically significant at the one percent level. Results using the rank list obtained in the survey from the subvillage head (r_e) are virtually identical to those from the community (r_c) (significant at the one percent level). This provides further evidence that the community and the elite broadly share the same assessments of welfare.

Perhaps most importantly, we find that the community treatment increases the rank correlation between the targeting outcomes and the self-assessments of own

TABLE 10—DO COMMUNITY MEETINGS REFLECT BROADLY SHARED PREFERENCES?

| | Meeting attendance (meeting data) (1) | Meeting attendance (HH data) (2) | Female attendance (meeting data) (3) | Mistarget (4) | Rank correlations with: | | | |
					Consumption (5)	Community survey (excl. subvillage head) (6)	Subvillage head survey (7)	Self-assessment (8)
Community treatment		0.349*** (0.042)		0.028 (0.021)	−0.089** (0.045)	0.232*** (0.040)	0.180*** (0.052)	0.072 (0.044)
Hybrid treatment	0.020 (0.029)	0.353*** (0.041)	0.008 (0.017)	0.026 (0.021)	−0.089** (0.044)	0.130*** (0.039)	0.064 (0.051)	0.046 (0.044)
Day meeting treatment	−0.021 (0.029)	0.013 (0.033)	0.104*** (0.017)	0.008 (0.016)	0.019 (0.033)	0.004 (0.029)	0.055 (0.038)	0.014 (0.033)
Elite treatment	−0.064** (0.029)	−0.300*** (0.033)	−0.085*** (0.017)	0.005 (0.016)	−0.004 (0.033)	−0.023 (0.029)	0.034 (0.038)	−0.017 (0.033)
10 Poorest treatment	0.022 (0.029)	0.023 (0.034)	−0.010 (0.018)	−0.006 (0.016)	0.031 (0.033)	0.047 (0.029)	0.044 (0.038)	0.062* (0.032)
Observations	431	287	428	5,753	640	640	640	637
Mean in PMT treatment		0.110		0.300	0.451	0.506	0.456	0.343

Note: For column 3, the dependent variable is the percentage of households in the village in which a female attends the meeting, using data collected from the meeting attendance lists.

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

poverty (r_s) by 10.2 percentage points, or about 30 percent of the level in the PMT (significant at 1 percent). The hybrid treatment also does so by 7.5 percentage points. The community targeting methods are thus more likely to conform with individual's self-identified welfare status.

B. *Are These Preferences Broadly Shared?*

The results above suggest that the ranking exercise moves the targeting process towards a welfare metric identified by community members. An important question is the degree to which this reflects the view of one group within the community. One experimental subtreatment was designed precisely to get at this question. In Table 10, we report the effect of changing the composition of the meeting by holding the meeting during the day, when women are more likely to be able to attend. We also consider the other subtreatments (elite and 10 poorest) in this analysis, as they could also plausibly have affected the welfare weights of those at the meeting.

We begin by investigating the impact of having a daytime meeting on attendance. This treatment does not change the share of households in the village that attend (columns 1 and 2). The percentage of households that are represented by women, however, is about 10 percentage points (for a total of 49 percent) higher in the day than during the evening meetings (column 3).

Although the day meeting treatment affected the gender composition of the meetings, columns 4–8 show that it did not affect the targeting outcomes. The elite treatment also did not affect the rank correlations with any of the various welfare metrics. Interestingly, the only subtreatment that affected the rank correlations was

the 10 poorest treatment, which increased the correlation of the treatments with ranks from self-assessments. Overall, there seems to be no evidence that the identity of the subgroup doing the ranking mattered.

VII. Understanding the Community's Maximand

The evidence so far suggests that the community has a systematic, broadly shared notion of welfare that is not based on per capita consumption, and that the community targeting methods reflect this different concept of welfare. This raises several key questions: Is the community simply mismeasuring consumption? Or does it value something other than consumption?

A. Does the Community Lack Information to Evaluate Consumption?

There is no definitive way to prove that the community has all the information that is available in the PMT.[37] The fact that those ranked early in the process were ranked at least as well as in the PMT suggests, however, that information is not the main constraint. We can also test whether the community has information about consumption beyond that in the PMT. To do so, we estimate

$$(4) \qquad RANKIND_{ijvk} = \alpha + \beta_1 RANKCONSUMPTION_{ivk}$$

$$+ \beta_2 RANKPMTSCORE_{ivk} + \nu_j + \varepsilon_{ijvk},$$

where $RANKIND_{ijvk}$ is household j's rank of household i (all ranks are in percentiles), $RANKCONSUMPTION_{ivk}$ is the rank of household i's per capita consumption in village v, and $RANKPMTSCORE_{ivk}$ is the rank of household i's PMT score that is computed using the baseline data. Fixed effects for the individual providing the ranking are included (ν_j), and standard errors are clustered at the village level. The results of this analysis are presented in column 1 of Table 11. In column 2, we instead include all of the variables that enter the PMT score separately rather than including the rank of the PMT score.

Table 11 illustrates that the community has residual information. Consumption is still highly correlated with individuals' ranks of other households from the baseline survey *even* after we control for the rank from the PMT (e.g., all the information that is contained in the PMT): a one percentile increase in consumption rank is associated with a 0.132 percentile increase in individual household ranks of the community (column 1). This is significant at the one percent level. In the more flexible specification (column 2), the correlation between consumption rank and survey rank remains positive (0.088) and significant at the one percent level.

The findings in Table 11 suggest that the community has residual information about consumption beyond that contained in the PMT score or even in the PMT variables.

---

[37] The ideal way to test this would be to ask households in the baseline survey to answer the questions in the PMT formula for other households in their village. This would be intrusive to do in a baseline survey, however, as it may make households feel as if they are "reporting" on other households in their village.

TABLE 11—INFORMATION

| | Community survey rank (r_c) | | | Survey rank |
	(1)	(2)		(2 continued)
Rank per capita consumption within village in percentiles	0.132*** (0.014)	0.088*** (0.012)		
Rank per capita consumption from PMT within village in percentiles	0.368*** (0.014)			
Household floor area per capita		0.001*** 0.000	Has this household ever got credit?	0.027** (0.011)
Not earth floor		0.060*** (0.010)	Number of children 0–4	0.000 (0.006)
Brick or cement wall		0.065*** (0.007)	Number of children in elementary school	0.003 (0.005)
Private toilet		0.047*** (0.008)	Number of children in junior high school	0.007 (0.007)
Clean drinking water		0.008 (0.009)	Number of children in senior high school	0.022*** (0.008)
PLN electricity		0.064*** (0.008)	Highest education attainment within HH is elementary school	0.007 (0.016)
Concrete or corrugated roof		0.027* (0.014)	Highest education attainment within HH is junior school	0.01 (0.016)
Cooks with firewood		0.031*** (0.008)	Highest education attainment within HH is senior high or higher	0.051*** (0.017)
Own house privately		0.034*** (0.008)	Total dependency ratio	0.004 (0.006)
Household size		0.004 (0.006)	AC	0.049** (0.023)
Household size squared		−0.001 (0.001)	Computer	0.045*** (0.011)
Age of head of household		0.011*** (0.002)	Radio/cassette player	0.001 (0.006)
Age of head of household squared		−0.000*** 0.000	TV	0.043*** (0.010)
Head of household is male		0.047** (0.019)	DVD/VCD player	0.017** (0.007)
Head of household is married		0.119*** (0.022)	Satellite dish	0.021* (0.011)
Head of household is male and married		−0.043* (0.026)	Gas burner	0.030*** (0.008)
Head of household works in agriculture sector		−0.006 (0.041)	Refrigerator	0.069*** (0.008)
Head of household works in industry sector		−0.043 (0.042)	Bicycle	−0.004 (0.007)
Head of household works in service sector		−0.018 (0.042)	Motorcycle	0.078*** (0.007)
Head of household works in formal sector		0.071 (0.045)	Car/minibus/truck	0.116*** (0.012)
Head of household works in informal sector		0.048 (0.045)	HP	0.014* (0.007)
Education attainment of HH head is elementary school		0.008 (0.008)	Jewelry	0.034*** (0.006)
Education attainment of HH head is junior school		0.036*** (0.010)	Chicken	−0.001 (0.006)
Education attainment of HH head is senior high school or higher		0.041*** (0.011)	Caribou/cow	0.065*** (0.012)
Observations		40,398		38,336

***Significant at the 1 percent level.
**Significant at the 5 percent level.
*Significant at the 10 percent level.

TABLE 12—WHAT IS THE COMMUNITY MAXIMIZING?

	Rank according to welfare metric			Targeting rank list in		
	Community survey ranks (r_c) (1)	Subvillage head survey ranks (r_e) (2)	Self-assessment (r_s) (3)	PMT villages (4)	Community villages (5)	Hybrid villages (6)
Log per capita consumption	0.176*** (0.008)	0.145*** (0.008)	0.087*** (0.004)	0.132*** (0.013)	0.197*** (0.014)	0.162*** (0.014)
Panel A. Household demographics						
Log HH size	0.164*** (0.011)	0.134*** (0.010)	0.073*** (0.006)	−0.028 (0.019)	0.154*** (0.019)	0.078*** (0.021)
Share kids	−0.125*** (0.021)	−0.094*** (0.021)	−0.037*** (0.012)	−0.296*** (0.035)	−0.068* (0.041)	−0.141*** (0.039)
Panel B. Ability to smooth shocks						
Elite connected	0.092*** (0.008)	0.044*** (0.009)	0.025*** (0.005)	0.062*** (0.016)	0.051*** (0.015)	0.043*** (0.015)
Total connectedness	−0.039*** (0.010)	−0.021** (0.009)	−0.015*** (0.005)	−0.016 (0.017)	−0.019 (0.017)	−0.054*** (0.019)
Number of family members outside subvillage	0.012*** (0.004)	0.010*** (0.003)	0.006*** (0.002)	0.020*** (0.006)	0.001 (0.006)	0.001 (0.006)
Participation through work to community projects	0.002 (0.011)	0.021** (0.010)	0.005 (0.006)	0.000 (0.018)	0.010 (0.019)	0.003 (0.019)
Participation through money to community projects	0.061*** (0.009)	0.041*** (0.009)	0.024*** (0.005)	0.056*** (0.016)	0.058*** (0.016)	0.034* (0.018)
Participation in religious groups	0.027*** (0.010)	0.033*** (0.010)	0.014** (0.006)	0.033** (0.016)	0.012 (0.017)	0.029 (0.017)
Total savings	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Share of savings in a bank	0.096*** (0.011)	0.069*** (0.010)	0.052*** (0.006)	0.121*** (0.018)	0.103*** (0.021)	0.075*** (0.020)
Debt as share of consumption	0.005*** (0.001)	0.004*** (0.001)	0.002*** (0.000)	0.002 (0.002)	0.007*** (0.001)	0.008*** (0.001)
Panel C. Discrimination against minorities?						
Ethnic minority	−0.024* (0.014)	−0.019 (0.014)	−0.003 (0.008)	0.012 (0.026)	−0.051** (0.025)	−0.011 (0.024)
Religious minority	0.012 (0.018)	−0.007 (0.017)	−0.014* (0.008)	−0.018 (0.030)	0.025 (0.032)	0.012 (0.033)
Panel D. Correcting for earnings ability						
HH head with primary education or less	−0.028*** (0.009)	−0.025*** (0.009)	−0.037*** (0.005)	−0.108*** (0.017)	−0.011 (0.018)	−0.066*** (0.017)
Widow	−0.104*** (0.014)	−0.083*** (0.014)	−0.012 (0.008)	0.009 (0.027)	−0.108*** (0.024)	−0.026 (0.028)
Disability	−0.045*** (0.016)	−0.037*** (0.014)	−0.026*** (0.008)	−0.079*** (0.027)	0.009 (0.026)	0.012 (0.027)
Death	−0.041* (0.025)	−0.031 (0.025)	−0.010 (0.015)	−0.111*** (0.042)	−0.013 (0.048)	−0.059 (0.043)
Sick	−0.038*** (0.011)	−0.041*** (0.011)	−0.028*** (0.006)	0.007 (0.018)	−0.018 (0.019)	−0.044** (0.019)
Recent shock to income	−0.001 (0.009)	−0.005 (0.009)	−0.013** (0.005)	−0.019 (0.016)	0.009 (0.016)	−0.012 (0.017)
Tobacco and alcohol consumption	−0.0002*** (0.000)	−0.0002*** (0.000)	−0.0001*** (0.000)	−0.0002*** (0.000)	−0.0002*** (0.000)	−0.0001*** (0.000)
Observations	5,337	4,680	5,724	1,814	1,876	1,889

Notes: Note that the children and household head education variables are explicitly included in the PMT regression (see Table 11). The PMT regression also includes dummies for the household head being male, married, and male × married, which together will be closely correlated with the widow variable.

*** Significant at the 1 percent level.
** Significant at the 5 percent level.
* Significant at the 10 percent level.

Moreover, the fact that almost all the PMT variables enter into the community ranks with plausible magnitudes suggests that the community has most of the information in the PMT as well, but chooses to aggregate it in different ways. While we cannot completely rule out that the community lacks some information that is present in the PMT, the evidence here suggests that differences in information are not the primary drivers of the different results.

B. *A Different View of Individual Welfare*

Table 12 explores the relationship between the welfare metrics (community survey rank r_c, elite survey rank r_e, and self-assessment rank r_s), the targeting results in PMT, community, and hybrid villages, and a variety of household characteristics that might plausibly affect either the welfare functions or the social welfare weights used in targeting. In columns 1–3, we present results of specifications where the dependent variable is the within-village rank of each household in the baseline survey according to different survey-based welfare metrics. In columns 4–6, the dependent variable is the treatment rank, put on a corresponding metric where the lowest-ranked (poorest) household in the dataset in each village is ranked 0 and the highest-ranked (richest) household in the dataset in each village is ranked 1.[38] We control for the log of per capita consumption in all regressions, and therefore the coefficients can be interpreted as conditional on per capita consumption. Thus, we identify where the community rankings deviate from a ranking based on consumption. We examine four dimensions on which villages may deviate: household demographic composition, ability to smooth shocks, discrimination against minorities or other marginal community members, and earning ability.

First, in panel A, we examine whether villages make adjustments for equivalence scales. In our setting, the PMT is explicitly defined using per capita consumption. Thus, it makes no adjustment for economies of scale in the household. By contrast, all of the community welfare functions (columns 1–3) reveal that people believe that there are household economies of scale, so that conditional on per capita consumption, those in larger households are considered to have higher welfare (as in Olken 2005). Likewise, the same is true for the community ranking, which assigns almost an identical household size premium (column 5). Interestingly, for a given household size and consumption, all methods rank households with more kids as poorer, even though children generally cost less than adults (Deaton 1997).

Second, the community may know more about other households' ability to smooth shocks. Conditional on current consumption, the household that is better able to smooth shocks may be at a higher long-run expected utility level and therefore may need transfers less. For example, if two families have the same per capita consumption, the one that is more elite-connected may worry less about bad shocks because it can expect to get help from rich relatives. The community might therefore feel that elite-connected households are richer than their consumption indicates. Whether or not this is the correct theory, it aligns perfectly with what we find. In panel B, we

---

[38] Note that some of the variables included as explanatory variables—including household size, share of kids, household head education, and widowhood—were explicitly included in the PMT regression, which may explain why some of these variables are significant predictors of targeting in the PMT regressions.

show that the community survey ranks put about a nine percentage point premium on being elite-connected, and even the elite and self-assessed survey ranks place a 4.4 and 2.5 percentage point premium, respectively. The community treatment ranks place a 5.1 percentage point premium on elite connectedness.

Similarly, there appears to be a premium for being better connected to the financial system. While total savings does not affect the rank, households that have a greater share of savings in a bank are classified as richer in both the individual surveys (column 1–3) and the community meeting (column 5). Households with family outside the village (who can presumably send remittances) are also ranked as less poor in terms of individual ranks, subvillage head ranks, and the self-assessment, though not in the community meetings.

Third, in panel C we test for discrimination against minorities or other marginal community members. We find no evidence of this: ethnic minorities are more likely to be ranked as poor in the community treatment, suggesting perhaps that even extra care is paid to them in the interest of social harmony (column 5). In addition, we find no evidence of favoring families that are more engaged with the community. Contributing labor to village projects does not affect a family's status. Those who contribute money, however, are viewed as rich (columns 1–3), though they are also likely to be ranked as richer by the PMT (column 4).

Finally, in panel D we find suggestive evidence that communities may try to provide the "right" incentives to households. For example, in a standard Mirrlees (1971) framework, one would ideally like to target on ability to earn, rather than actual earnings, so as not to disincentivize households. To test whether this is what communities are doing in practice, we first look at the education level of the household head. Households where the household head has a primary education or less rank two to four percentage points poorer, conditional on their actual consumption. Similarly, households headed by a widow, those with a disability, and those where there is a serious illness are all rated poorer, conditional on actual consumption. The adjustment for widowhood is also reflected in the community treatment ranking, but not the disability and serious illness adjustments (column 5).[39] Finally, and rather interestingly, the village does not penalize those who spend a lot of money on smoking and drinking. Families with these attributes are actually ranked lower both in community surveys (columns 1–3) and community meetings (column 5), suggesting that the village treats these preferences as problems for the family as a whole rather than as behaviors that should be punished.

VIII. Conclusion

The debate regarding decentralization in targeting is usually framed in terms of the benefits of utilizing local information versus the costs of some form of malfeasance, such as elite capture. While we started with an experiment that took both of these ideas very seriously, our results point to a third factor as being very important: the

---

[39]There are, of course, two interpretations of these findings. One interpretation is that households are conditioning on earnings ability—i.e., if you are highly educated but do not earn much, that is your fault and you should not receive subsidies for it. Another interpretation, however, is that education is merely another signal of poverty that is more easily observable to the community than actual consumption, though communities would need to be over-weighting this signal for this effect to produce a negative coefficient conditional on actual consumption.

community seems to have a widely shared objective function other than per capita consumption, and implementing this objective is a source of widespread satisfaction in the community. Moreover, this objective function does not differ based on elite capture. Rather, these preferences appear to be informed by a better understanding of factors that affect a household's earning potential or vulnerability, such as the returns to scale within the family, as compared to relying purely on consumption.

Given these findings, if targeting the poor based on consumption is the only objective, the PMT does perform somewhat better than the community methods, though the difference between the methods in terms of the ultimate poverty impact for a typically sized program is not significantly different. Especially given the relatively small differences in ultimate poverty outcomes between the alternative treatments, it is not evident that there is a strong enough case to overrule the community's preferences in favor of the traditional consumption metric of poverty, especially given the gain in satisfaction and legitimacy. On the other hand, what is clear is that there is no case for the intermediate hybrid method: it resulted both in worse targeting performance and low legitimacy. This may be because its main theoretical advantage—preventing elite capture—was not important in our setting. It is possible that perhaps alternative hybrid designs (e.g., using a PMT process in the first stage but then allowing the community to add some very poor households to the resulting beneficiary list) might perform better than those where the selection process is ultimately determined strictly by the PMT survey results, as the community does better at identifying the very poor.

The findings in this paper raise several interesting questions. First, while we found little evidence of elite capture, it is possible that this might change over time as individuals learn to better manipulate the system. Manipulation over time has been shown to occur in some kinds of PMT systems (Camacho and Conover 2011), but whether it would occur when the per-village allocation is fixed, and whether it would be more or less severe in community-targeted systems, are important open questions. Second, given how well the community outcomes match individual self-assessments, an important question is whether a self-targeting system (perhaps connected to an ordeal mechanism, as in Nichols and Zeckhauser 1982) could provide a more cost-effective method. We regard these as important questions for future research.

REFERENCES

Alatas, Vivi, Abhijit Banerjee, Rema Hanna, Benjamin A. Olken, and Julia Tobias. 2012. "Targeting the Poor: Evidence from a Field Experiment in Indonesia: Dataset." *American Economic Review*. http://dx.doi.org/10.1257/aer.102.4.1206.

Alderman, Harold. 2002. "Do Local Officials Know Something We Don't? Decentralization of Targeted Transfers in Albania." *Journal of Public Economics* 83 (3): 375–404.

Bardhan, Pranab, and Dilip Mookherjee. 2005. "Decentralizing Antipoverty Program Delivery in Developing Countries." *Journal of Public Economics* 89 (4): 675–704.

Camacho, Adriana, and Emily Conover. 2011. "Manipulation of Social Program Eligibility." *American Economic Journal: Economic Policy* 3 (2): 41–65.

Cameron, Lisa A. 2002. "Did Social Safety Net Scholarships Reduce Drop-Out Rates during the Indonesian Economic Crisis?" The World Bank, Policy Research Working Paper 2800.

Coady, David, Margaret Grosh, and John Hoddinott. 2004. "Targeting Outcomes Redux." *World Bank Research Observer* 19 (1): 61–85.

Conn, Katharine, Esther Duflo, Pascaline Dupas, Michael Kremer, and Owen Ozier. 2008. "Bursary Targeting Strategies: Which Method(s) Most Effectively Identify the Poorest Primary School Students for Secondary School Bursaries?" Unpublished.

Daly, Anne, and George Fane. 2002. "Anti-Poverty Programs in Indonesia." *Bulletin of Indonesian Economic Studies* 38 (3): 309–29.

Deaton, Angus. 1997. *The Analysis of Household Surveys: A Microeconometric Approach to Development Policy*. Baltimore: Johns Hopkins University Press for the World Bank.

Galasso, Emanuela, and Martin Ravallion. 2005. "Decentralized Targeting of an Antipoverty Program." *Journal of Public Economics* 89 (4): 705–27.

Mirrlees, James A. 1971. "An Exploration in the Theory of Optimum Income Taxation." *Review of Economic Studies* 38 (114): 175–208.

Nichols, Albert L., and Richard J. Zeckhauser. 1982. "Targeting Transfers through Restrictions on Recipients." *American Economic Review* 72 (2): 372–77.

Olken, Benjamin A. 2005. "Revealed Community Equivalence Scales." *Journal of Public Economics* 89 (2–3): 545–66.

Olken, Benjamin A. 2006. "Corruption and the Costs of Redistribution: Micro Evidence from Indonesia." *Journal of Public Economics* 90 (4-5): 853–70.

Olken, Benjamin A. 2010. "Direct Democracy and Local Public Goods: Evidence from a Field Experiment in Indonesia." *American Political Science Review* 104 (2): 243–67.

Ravallion, Martin. 2008. "Miss-targeted or Miss-measured?" *Economics Letters* 100 (1): 9–12.

Ravallion, Martin. 2009. "How Relevant Is Targeting to the Success of an Antipoverty Program?" *World Bank Research Observer* 24 (2): 205–31.

Ravallion, Martin, and Lorraine Dearden. 1988. "Social Security in a "Moral Economy": An Empirical Analysis for Java." *Review of Economics and Statistics* 70 (1): 36–44.

Seabright, Paul. 1996. "Accountability and Decentralisation in Government: An Incomplete Contracts Model." *European Economic Review* 40 (1): 61–89.

Survei Sosial Ekonomi Nasional (**National Socio-Economic Survey—Susenas**)**.** 2007. Badan Pusat Statistik (Indonesian Bureau of Statistics—BPS).

World Bank. 2008. *Global Purchasing Power Parities and Real Expenditures: 2005 International Comparison Program*. Washington, DC: World Bank.

World Bank. 2006. "Making the New Indonesia Work for the Poor: Overview." Washington, DC: World Bank.

World Bank Urban Poverty Project Data. 2007. World Bank Jakarta.