Corruption and the costs of redistribution: Micro evidence from Indonesia

Benjamin A. Olken

National Bureau of Economic Research, USA

Abstract

This paper examines the degree to which the corruption in developing countries may impair the ability of governments to redistribute wealth among their citizens. Specifically, I examine a large anti-poverty program in Indonesia that distributed subsidized rice to poor households. I estimate the extent of corruption in the program by comparing administrative data on the amount of rice distributed with survey data on the amount actually received by households. The central estimates suggest that, on average, at least 18% of the rice appears to have disappeared. Ethnically heterogeneous and sparsely populated areas are more likely to be missing rice. Using conservative assumptions for the marginal cost of public funds, I estimate that the welfare losses from this corruption may have been large enough to offset the potential welfare gains from the redistributive intent of the program. These findings suggest that corruption may impose substantial limitations on developing countries' redistributive efforts, and may help explain the low level of transfer programs in developing countries. © 2005 Elsevier B.V. All rights reserved.

Keywords: Corruption; Transfer programs

1. Introduction

Transfer programs and social safety nets have the potential to substantially improve welfare in the developing world. However, the large amount of funds involved in transfer programs may be an enticing target for potentially corrupt officials. In many developing countries, poor financial development means that implementing a transfer program often involves physically moving cash or in-kind benefits through multiple layers of a bureaucracy,

in which case opportunities for corruption may be particularly rife. If the losses due to corruption are large enough, they may outweigh the welfare benefits from redistribution.

This paper explores empirically the degree to which corruption may inhibit redistribution by examining the extent of corruption in a large Indonesian transfer program that distributed heavily subsidized rice to poor households. Studying corruption empirically is inherently difficult, as corrupt officials go to substantial lengths to conceal their activities. I measure the extent of corruption in the Indonesian transfer program by obtaining government administrative records on the amount of subsidized rice allocated to particular villages and districts, and comparing these records to household surveys that ask whether villagers actually received the rice. I perform this analysis using two separate data sets, one a detailed survey of 100 villages and one a nationally representative survey of over 200 Indonesian districts.

Using this approach, I find that at least 18% of the program's rice disappeared between the time it left government warehouses and the time it reached households. In many respects, this estimate constitutes a lower bound, and as discussed in the text, the actual amount of corruption may be substantially higher. In both datasets, the missing rice appears to be highly concentrated in a small number of areas—for example, in the 100 villages dataset, almost 60% of the detected missing rice comes from only 10% of villages. This concentration suggests that the missing rice is not driven by systematic underreporting by villagers—if it was, the distribution of the missing rice would likely be much more uniform.

I then examine in which types of areas the missing rice is most prevalent. In both datasets, I find that ethnically fragmented areas are more likely to be missing rice. I also find that missing rice is more prevalent in sparsely populated areas, where monitoring may be more difficult. I also find weaker, suggestive evidence that missing rice is more prevalent in poorer areas, and in areas with fewer social organizations.

In the final section of the paper, I perform a welfare calculation to compare the costs of this corruption to the potential redistributive benefits from the program. I find that, even under conservative assumptions for the marginal cost of public funds, corruption was sufficiently large to turn an otherwise welfare-enhancing program into a program that may have been welfare-reducing on net.

Most empirical work on corruption to date has been dominated by the use of subjective assessments of corruption. (See Rose-Ackerman (2004) for a survey.) This paper, however, is related to a small but growing recent literature that seeks to measure theft of funds directly by comparing two measures of the same quantity, one "before" and one "after" the corruption takes place. For example, Reinikka and Svensson (2004) measure corruption in Uganda by comparing central government data on public grants to schools with a survey of school officials to determine what fraction of the grants were received. Olken (2005) measures corruption in Indonesian road projects by comparing villages' official expenditure reports on road construction with engineers' independent assessments of the roads' cost. Di Tella and Schargrodsky (2003), who examine corruption in hospital procurement, compare data on prices paid for hospital inputs with market prices.[1] Unlike

---

[1] Although not exploring theft of funds, the other context in which a similar 'before corruption' and 'after corruption' has been used is international trade, where several authors have measured tax evasion by comparing trade data from the exporting and importing country for the same bilateral trade flows. (Yang, 2003; Fisman and Wei, 2004).

each of these previous studies, which required the use of special-purpose surveys, this study provides, to the best of my knowledge, the first example of how corruption can be detected by combining central administrative data with the types of household surveys routinely carried out by statistical bureaus in many countries. The approach used here can thus be carried out on a country-wide scale, at very low additional cost, in a wide variety of countries and contexts.

The remainder of the paper is organized as follows. Section 2 discusses the background of the Indonesian poverty alleviation program I study, known as the *OPK* program. Section 3 presents the estimates of the extent of corruption in the program. Section 4 examines in which types of settings corruption appears to be most prevalent. Section 5 presents the welfare calculation used to compare the costs of corruption with the benefits from redistribution. Section 6 concludes.

2. Background: corruption and the *OPK* Program

In 1997–1998, Indonesia experienced a severe economic collapse. In response, the government of Indonesia implemented several new social safety net programs, the largest of which was the *Operasi Pasar Khusus* (Special Market Operation), or *OPK*. The program (renamed *Raskin* in 2001) is currently the largest redistributive program in Indonesia.

The OPK program provided income support in the form of subsidized rice, the principal staple in Indonesia. During the period I study, eligible households were allowed to purchase 20 kg of OPK rice per month, at a price approximately 60% below market. The size of this subsidy was substantial—for the median eligible household purchasing their full allotment of rice, the subsidy represented approximately 9% of total pre-program monthly household expenditures.[2] Approximately 50% of rural Indonesian households were eligible to receive the subsidized OPK rice.

The mechanism through which OPK rice was distributed provides some sense of how corruption in OPK may have occurred. Village governments were in charge of distributing OPK rice to households. Each month, the village's allotment of rice would be retrieved from the nearest government warehouse by village officials, usually the village head, or would be delivered to the village office by the government logistics agency. Village heads were responsible for dividing the rice, received in bulk, into 20 kg sacks for purchase by households, for designating which households could purchase the rice, and for collecting the copayment. So long as the central government received the copayment from the villages, there was virtually no monitoring by the central government of how the rice was distributed within the villages.

Fieldwork I conducted in 2001 in OPK villages suggested, at least anecdotally, that corruption may have been a serious problem in OPK. For example, in interviews in one village in August 2001, residents reported that the rice deliveries had become intermittent in 2000, and had stopped completely by the beginning of 2001. Yet, the government warehouse reported that 1.6 tons of rice had been distributed to the village on time every month. (See Olken et al., 2001.) Based on the available evidence, it seemed as though the village head,

---

[2] Author's calculations using the SUSENAS dataset.

possibly in collusion with other the village heads in the subdistrict, was intercepting the rice on its way from the warehouse to the village, and reselling it secretly on the private market. For the most part, villagers had no source of information regarding their village's rice allocation other than the village head, making it possible for the village head to siphon off rice on its way to the village. This lack of transparency, combined with poorly publicized changes in program rules, provided substantial opportunity for corruption.

3. Estimating the amount of "Missing Rice"

To estimate the amount of missing rice, one needs to estimate two quantities: the total amount of rice that was *disbursed* by the central government to an area and the total amount of rice that was *received* by villagers in that area. The difference between the amount disbursed by the government and the amount received by households represents an estimate for the amount that disappeared in the process. I begin by describing my approach to estimating each of these two quantities, and then present the results.

3.1. Estimating the amount of rice received by villagers

To estimate the amount of rice received by villagers, I use household level survey data from two datasets, the 1998–1999 Hundred Villages Survey, known in Indonesian as the *Survei Seratus Desa*, or SSD, and the 1999 National Social Welfare Survey, or SUSENAS. The SSD is a four-wave panel dataset conducted by the Indonesian Central Statistics Office in 100 poor communities throughout Indonesia. Each wave of the SSD covered 120 households in each village, and is representative at the village level. The SUSENAS is a national survey consisting of 200,000 respondents, and is representative for each of Indonesia's 223 districts.[3] Summary statistics from the SSD and the SUSENAS are presented in Table 1. They confirm that villages surveyed in the SSD are poorer, and more remote, than the national average.

The final two waves of the SSD contain information on whether each household received OPK rice in the three or four months prior to the survey.[4] As shown in Table 1, 56% of SSD households report receiving OPK rice. Since the SSD is representative at the village level, one can use this question to obtain an estimate of the percentage of households in each of the 100 villages in the SSD receiving OPK rice in the period prior to each wave of the survey. Similarly, the 1999 SUSENAS asked households how many times they had

---

[3] In addition to the 223 districts, the SUSENAS also collects data on Indonesia's 64 *kotamadya*, or major cities. I exclude these major cities from the analysis, as there were other subsidized rice programs in major cities.

[4] Technically, the question does not specifically refer to OPK; instead, it asks if households received "free or subsidized *sembako*." *Sembako* refers to the nine basic staple commodities, of which rice accounts for by far the largest share of expenditures. However, as OPK was the only large-scale *sembako* program in rural areas, I follow other authors (e.g., Pritchett et al., 2002) in interpreting this question as referring primarily to OPK. To the extent that households also reported non-OPK assistance as *sembako* assistance, this approach will underestimate the amount of corruption in the program. The second wave of the SSD also contains a question on OPK receipt, but since the OPK program was still being phased in at the time the that wave of the survey was fielded, I exclude this wave from the analysis.

Table 1
Summary statistics from the SSD and SUSENAS

	SSD	SUSENAS
Household characteristics		
Log HH expenditure per capita	11.43	11.71
	(0.47)	(0.40)
HH size	4.17	4.09
	(1.78)	(1.77)
Does HH have electricity	0.591	0.812
	(0.492)	(0.390)
Did HH receive OPK rice?	0.561	0.404
	(0.496)	(0.491)
Was HH eligible to receive OPK rice?	0.668	0.533
	(0.471)	(0.498)
Village characteristics		
Village population	3655	5458
	(2568)	(5192)
Share of population classified as poor (KPS or KS1)	0.588	0.464
	(0.263)	(0.268)
Distance to nearest district capital (km)	77.12	57.89
	(74.03)	(64.14)

Notes: Standard deviations in parentheses. SUSENAS sample excludes major cities (*kotamadya*). For comparison with the SUSENAS, which was fielded in January 1999, SSD expenditure data is from Wave II (December 98); all other SSD household characteristics are from Waves III (May 1999) and IV (October 1999).

received subsidized rice in the 6 months prior to the survey, which was conducted in January 1999. Since the SUSENAS is representative at the district level, one can use this data to estimate the percentage of households in each district that received the rice.

There are, however, several issues with using these surveys to estimate the amount of rice received by villagers. First, the SSD only provides information on *whether* the individual receives rice, and the SUSENAS only provides information on the *number of times* a household received rice. Neither provides information on the *quantity* of rice received. I therefore make, as a baseline assumption, the assumption that each household that received rice received the full official monthly allotment of rice.[5]

Quantitative and qualitative evidence suggests that this assumption is quite generous, and therefore likely to produce an underestimate of the amount of corruption. In particular, a 1999 Indonesian survey found that of the households receiving subsidized rice, only 19% of households received the full 20 kg, and in fact 68% of households received less than *ten* kg of rice.[6] In addition, qualitative assessments of OPK have tended to find that

---

[5] The official allotment was raised from 10 to 20 kg/month in December 1998. For the SSD waves, both of which occurred after this time, I therefore assume that each household received 20 kg. For the SUSENAS, I assume that the first time they received rice (e.g., December 1998), they received 20 kg; each additional time they received rice (e.g., prior to December 1998) I assume they received 10 kg.

[6] The survey was conduction by an Indonesian think tank, the Economic and Social Research, Education, and Information Institute, known in Indonesian as LP3ES. For details, see LP3ES, 2000. Essentially no households report receiving *more* than the official 20 kg allotment; the less than 1% of households that did were all located in one remote province not included in the SSD.

households received substantially less than 20 kg of rice, and never found evidence of any households receiving more (see e.g., Olken et al., 2001). In the results below, I therefore also report estimates of corruption based on less conservative assumptions about the amount of rice received.

The second issue with using the SSD is that it asks if households received rice in the previous three or four months (depending on the wave), while the OPK program distributed rice monthly. To the extent that the set of households that received rice varied from month to month, more households will report receiving rice in the previous three or four months than received it in any given month. Failing to correct for this would overstate the amount of rice received and understate the amount of corruption in the program. I use the panel aspect of the OPK data to estimate the degree to which the recipient list varied from month to month, which I use to correct for this problem. Details about this correction can be found in Appendix A.

3.2. Estimating the amount of rice disbursed to each village

For information on the amount of rice that each village *should have* received, I use administrative records on rice distribution. First, information on the amount distributed to individual villages was kept at the district level distribution centers. Accordingly, for each of the 100 villages in the SSD, I obtained from each district level distribution center the number of kilograms of rice distributed each month to each village in the SSD. Second, for the district-level data from the SUSENAS, I obtained from the central government information on the total amount of rice distributed in each district each month. (See Tabor and Sawit, 1999).

As discussed above, the household data in the SSD only asked households whether they had received rice, or in the case of the SUSENAS, how many times they had received rice. I use these survey data to compute how much rice was actually distributed, under the generous assumption that each household that received rice received the maximum amount. I then take the total amount I can account for in the village/district, and compare it to the amount of rice distributed to the village/district during the same period of time. Only when the maximum amount of rice received by households is less than the amount disbursed to the village/district do I conclude that rice was missing from that village/district.

3.3. Results

The results from this procedure are presented in Table 2. The rows represent different assumptions about the amount of rice households received, as a percentage of the official allocation. The first column presents results from the SSD, where the data are aggregated to the village level. Under the assumption that each household receiving rice received the full 20 kg, i.e., 100% of the official allocation, I estimate that 17.8% of the rice appears to be missing. This estimate, based on the finest level of detail available, is the estimate I consider the central estimate of the paper. However, even this 17.8% estimate may still be substantially below the actual amount of missing rice. If, for example, the maximum amount households were assumed to receive was 75 of the official allocation (i.e., 15 kg instead of 20 kg per month), the estimate of the percentage rice missing jumps to 28.0%.

Table 2
Estimates of missing rice

| Assumption: percent of official allocation received by each household	(1) Village level aggregation (SSD data)	(2) District level aggregation (SSD data)	(3) District level aggregation (SUSENAS data)
75%	.280	.136	.232
100%	**.178**	.087	.093
125%	.120	.066	.052

Notes: Each cell represents the estimated share of rice unaccounted for, using data from the SSD (columns 1 and 2) and from the SUSENAS (Column 3). Column 1 presents data aggregated to the village level, columns 2 and 3 present data aggregated to the district (*kabupaten*) level. For SSD data, I assume that α, the fraction of households whose rice receipt is assumed to have alternated each month, is equal to 0.15, as implied by the panel aspect of the data (see Appendix 1). The central estimate of the paper, based on village-level aggregation from the SSD and assuming 20 kg of rice received by each household, is shown in bold.

An interesting question is whether all of this missing rice is coming from just a few villages, or whether the theft was spread more evenly across many villages. Fig. 1 shows the distribution of the missing rice from the SSD among village-waves (since there are two waves of data for each village, there are two corruption calculations for each village, and they are entered separately when creating this Figure). Specifically, it plots the percentage of the total amount of the missing rice attributable to each decile of village-waves. The figure shows that the missing rice was quite concentrated—all of the missing rice came from just 32% of village-waves, and 61% of the missing rice was concentrated in the top 10% of village-waves.

It is interesting to compare these estimates to an independent set of estimates from the nationally-representative SUSENAS. As discussed above, the SUSENAS is representative only at the district level, rather than the village level. Aggregating at the district level rather than the village level underestimates the percent of rice missing, as villages where the rice was distributed widely (and where there were many recipients, each receiving a small amount of rice) mask corruption in other villages in their district, since I assume that all households receiving rice received the full 20 kg allotment.

To compare the SSD estimates to the SUSENAS estimates, we need to first re-estimate the missing rice in the SSD dataset at the same level of aggregation as the SUSENAS. These estimates are presented in Column 2 of Table 2. As expected, when I aggregate to the district level using the same data from the SSD, the estimated rice missing falls from 17.8% to 8.7%. Column 3 then presents the district-level estimates from the SUSENAS. The results suggest that, if all households receiving rice received the full allocation, then 9.3% of the rice is missing. This is quite close to the 8.7% estimated at the same level of aggregation using the SSD. This suggests that, were village-level data similar to the SSD available nationally, the results would be very similar to the 17.8% missing rice found in the SSD.

Fig. 2 shows the concentration of the missing rice by decile of districts in the SUSENAS. The pattern of the missing rice is, if anything, even more concentrated than the village data from the SSD shown in Fig. 1, with 66% of missing rice concentrated in the top 10% of districts.
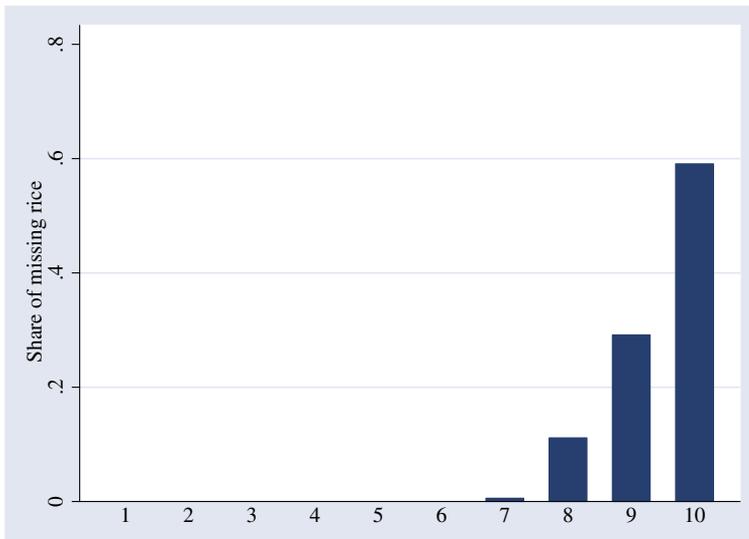
Notes: Data is from SSD. The unit of observation is a village-wave. Each bar reports the share of the total missing rice attributable to that decile of village-waves. Results using the village, rather than the village-wave, as the unit of observation are qualitatively similar.

Fig. 1. Distribution of missing rice across villages, by decile.
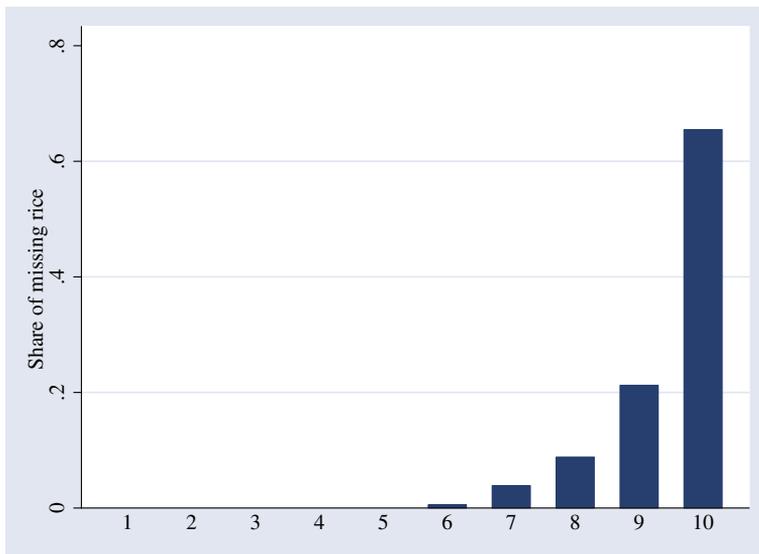
Notes: See Notes for Figure 2. Data is from SUSENAS, and the unit of observation is the district.

Fig. 2. Distribution of missing rice across districts, by decile.

I explored several potential concerns with the analysis. One such concern is that the estimates of missing rice are driven by underreporting of OPK receipt by households in the survey. Overall, however, reporting of OPK receipt is actually much *higher* than the amount of rice given out. For example, in the SSD, the official allocations of rice were such that only 33% of households should have received rice if each household received the official allocations. In fact, 56% of households report receiving OPK rice in the 3 months previous to the survey. As discussed above, the fact that the percent reporting rice receipt exceeds the amount distributed is likely because some villages distributed the rice more widely than intended (with each recipient receiving a smaller amount), and due to churning of households on the recipient list from month to month. Nevertheless, this high level of reporting suggests that underreporting systematic enough to explain the missing rice is unlikely to be a major issue.

The concentration of the missing rice provides even stronger evidence that under-reporting is unlikely to be driving the estimates of the missing rice. To see this, note that if the probability a household failed to report rice receipt was relatively constant across villages, and if this was driving the estimates of corruption, the distribution of missing rice would be relatively uniform across villages, and all of the missing rice would be coming from a large number of villages each missing only a small fraction of their rice allocation. In fact, however, most of the missing rice comes from small number villages missing a large percentage of the rice, which suggests that this type of survey underreporting is unlikely to be driving the results.

One can develop statistical tests along these lines. As discussed in Appendix B, I use Monte-Carlo simulations to develop the expected distribution of missing rice in the SSD, conditional on the missing rice in fact being driven by under-reporting. I consider both underreporting that is uniform—i.e., everyone's probability of receiving rice is reduced by the 17.8% missing I observe in the actual data—and a case where underreporting is correlated with observable household characteristics. All of the scenarios I consider produce distributions much more uniform than the extreme concentration of missing rice found in the data and shown in Fig. 1. All of the tests reject the null hypothesis that the missing rice is driven by under-reporting with p-values less than 0.01.

One final possibility is that the dramatic amount of missing rice in certain areas may be driven by bad data collection in these villages or regions. To examine this, I examined several other variables unrelated to corruption. Specifically, I computed the mean number of years of education and mean age in the population in the SSD/SUSENAS, and compared them to an independent estimate from the 2000 Population Census. I use the absolute value of the difference in logs between the SSD/SUSENAS estimates and the corresponding Census estimates as a measure of how much the SSD/SUSENAS disagrees with the Census data in a given geographic area, and hence as a measure of possible data quality problems in that area. In the SSD data, which are the data used for the central estimates of the paper, I find no relationship between these 'data mismatch' measures and the probability of my detecting missing rice in a village. In the district-level data in the SUSENAS, these variables do have some predictive power for detecting missing rice in the district, though conditional on a district having any missing rice, these 'data mismatch' measures do not predict the amount of rice missing. All told, the extreme concentration of

the missing rice, combined with the evidence on data quality, suggest the missing rice is unlikely to be driven primarily by data reporting issues.

4. Where does corruption occur?

A natural next step is to investigate the correlates of corruption—i.e., which social, economic, and political factors appear to be related to higher levels of corruption. For example, the cross-country literature, using data on perceived levels of corruption, has suggested that there is more corruption in countries that are poorer, more ethnically fragmented, and of French or Socialist legal origin (Mauro, 1995; LaPorta et al., 1999; Treisman, 2000). Looking across U.S. states, Glaeser and Saks (2004), using data on Federal convictions for corruption, also find that poorer states and more racially heterogeneous states, as well as states with greater inequality, experience higher levels of corruption. However, there is much less within-country evidence on the determinants of corruption in developing countries, particularly using objective, rather than subjective, measures of corruption.[7]

To examine these hypotheses, I estimate the relationship between detecting missing rice in an area and a range of characteristics about the area. I estimate this relationship both at the village level using SSD data and at the district level using SUSENAS data. The fact that I estimate these regressions uses two separate sources of data, at two different levels of aggregation, lends further credibility to the results.

I examine the relationship between area characteristics and corruption by estimating the following Probit specification:[8]

$$\text{Prob}\big(\text{Missing Rice}_{ij} = 1\big) = \Phi\big(\alpha_j + \beta * X_{ij}\big) \tag{1}$$

where i represents the unit of observation (village in SSD data, district in SSD data), X_{ij} are a set of village/district characteristics, α_j are a set of dummies for each of Indonesia's six major island groups, and Φ is the Normal CDF.[9] The independent variables X_{ij} include two moments of the per-capita expenditure distribution—the median level of per-capita expenditure in the village/district and the ratio of the 90th percentile of per-capita expenditure in the village/district to the 10th percentile of per-capita expenditure. These are calculated directly from the SSD and SUSENAS surveys. They also include a number of village characteristics, such as village population density, the presence of roads usable year-round, the number of types of social organizations in the village. In some specifications (columns 2 and 4 of Table 3), I also include the percent of households officially listed as 'poor' (pre-prosperous (KPS) or prosperous level 1 (KS1)), which was

---

[7] One notable exception is Reinikka and Svensson (2004), who in their study of Uganda find more corruption of school aid in poorer communities.

[8] Estimation using a Tobit specification, where the dependent variable is the percentage of rice missing (truncated at 0), produce similar qualitative results to the Probit results presented here. Similar results are also obtained using linear probability models, or using a Probit with the dependent variable a dummy that takes a value of 1 if more than 50% of the rice in the village/district was missing.

[9] The island groups are Java, Sumatra, Kalimantan, Sulawesi, Nusa Tenggara (including Bali), and Maluku. Irian Jaya is dropped from the analysis, and so it's not included here.

Table 3
Determinants of missing rice

| | (1) | (2) | (3) | (4) |
	SSD	SSD	SUSENAS	SUSENAS
Median log per. cap. expenditure	−0.283	−0.355*	−0.135	−0.020
	(1.61)	(1.94)	(0.47)	(0.06)
Inequality (90–10 ratio)	−0.128	−0.131	0.112	0.097
	(1.32)	(1.33)	(1.12)	(0.96)
Population density (households per hectare)	−0.006*	−0.006*	−0.034*	−0.032*
	(1.93)	(1.95)	(1.93)	(1.85)
Year-round roads	0.056	0.044	0.596	0.598
	(0.40)	(0.31)	(1.29)	(1.31)
Mean years education	0.037	0.032	−0.011	0.008
	(0.88)	(0.77)	(0.19)	(0.12)
Village ethnic fragmentation	0.964***	1.005***	0.729*	0.739*
	(3.55)	(3.44)	(1.87)	(1.89)
Village religious fragmentation	0.115	0.153	−1.258**	−1.292**
	(0.33)	(0.44)	(2.36)	(2.43)
Number social organizations	0.007	0.005	−0.118***	−0.112***
	(0.29)	(0.21)	(3.57)	(3.18)
Share officially poor		−0.161		0.313
		(0.85)		(0.91)
Island group fixed effects	YES	YES	YES	YES
Observations	190	190	200	200
Mean dep. var	0.32	0.32	0.43	0.43

Notes: Robust z statistics in parentheses. Estimation is by Probit, and marginal effects are reported at the mean levels of the independent variables. Columns (1) and (2) report data from the SSD, where the dependent variable is a dummy equal to 1 if there is missing rice in the village. There are two observations per village, one for each wave used in the analysis, and standard errors are adjusted for clustering at the village level. Columns (3) and (4) report data from the SUSENAS, where the dependent variable is a dummy equal to 1 if there is missing rice in the district. In columns (3) and (4), population-weighted averages are used for village characteristics. Percent officially poor was the variable the central government used for determining the amount of rice allocated to each geographic area.
* Significant at 10%; ** significant at 5%; *** significant at 1%.

the criteria used by the government in determining how much rice to allocate to a given area.[10] These characteristics are obtained from the 1999 Census of Villages. I also include the mean years of education, as well as Herfindahl index of village ethnic fragmentation and village religious fragmentation, using data from the 2000 Population Census. For the village variables in the district level regressions, I take the population-weighted average levels of these characteristics for all villages in each district.

It is important to note that the coefficients in Eq. (1) should be interpreted with some caution. To see this, note that some villages stole rice, some distributed small amounts of rice to many more households than those eligible, and others may have done both. Because I assume that each household that received rice received the maximum amount possible, I will only detect rice missing rice in those villages that both stole rice *and* did not

---

[10] I do not include the actual amount of rice allocated as a regressor, as measurement error in that variable will introduce a mechanical correlation between the amount of rice allocated and the presence of missing rice.

substantially increase the number of households receiving rice. The coefficients in Eq. (1) thus tell us which types of villages were likely to both steal rice and not to mask it by distributing a small amount of rice to many people.

The results are presented in Table 3. The coefficients presented are the marginal effects computed from estimating Eq. (1). Columns 1 and 2 present village-level results using SSD data, and columns 3 and 4 present district-level results using the SUSENAS.

Several results stand out. First, in both datasets I find that areas with higher within-village ethnic fragmentation have a higher likelihood of experiencing corruption, consistent with the cross-country evidence. Interestingly, the evidence on religious fragmentation is much less conclusive, with religious fragmentation having a negative effect in the SUSENAS and no effect in the SSD. Second, in both datasets, I find less corruption in areas with higher population density. This may be because there is less information, and therefore weaker monitoring, in sparser areas. Third, although not statistically significant in most specifications presented, there is weak evidence that richer areas have lower corruption, although there is no clear evidence on the impact on inequality. Finally, the results from the SUSENAS indicate that areas with more social organizations, such as community-self help groups, religious study groups, and women's organizations, are less likely to have missing rice, though these results are not found in the SSD data. These results, particularly on ethnic fragmentation and income levels, are broadly consistent with the results found in cross-country studies.

5. Welfare implications: were redistributive attempts on net welfare-reducing?

The paper thus far has presented evidence that corruption in the program was substantial—the central estimates suggest that 18% of the rice appears to have disappeared. This section performs a welfare calculation to compare the welfare losses caused by corruption to the potential welfare gains achieved by redistribution in the program. The welfare losses from this corruption entail both the foregone redistribution from the stolen rice, as well as the additional costs imposed by the dead-weight loss from the taxation used to pay for the missing rice. I ask whether, under reasonable parameters for the marginal cost of public funds, these welfare losses were large enough to make a program such as OPK welfare-reducing on net.

To do this, I assume a utilitarian social welfare function with CRRA individual utility and equal welfare weights. I present results for the case with the coefficient of relative risk aversion $\rho = 1$ (log utility) and $\rho = 2$. Note that as the coefficient of relative risk aversion increases, the social welfare function implicitly places more and more weight on lower income households, and therefore the welfare gain from a given amount of redistribution increases.

I use this social welfare function to calculate four different social welfare levels: 1) the social welfare actually achieved by the program, 2) the social welfare that would have been achieved with the same set of beneficiaries but in the absence of corruption, 3) the social welfare that would have been achieved had only eligible households received the rice (also in the absence of corruption), and 4) the social welfare in the absence of the program. For households receiving the rice in each scenario, I assume

Table 4
Comparing costs and benefits

Allocations:		Utilitarian, CRRA utility $\rho = 1$ (% of welfare maximizing utility)	Utilitarian, CRRA utility $\rho = 2$ (% of welfare maximizing utility)
Program	Actual allocation	52.23	35.31
	Actual allocation, no corruption	62.06	42.73
	Official eligibility guidelines	60.90	42.10
No program	Consumption tax, MCF = 1.00	46.90	24.68
	Consumption tax, MCF = 1.20	56.25	29.59
	Consumption tax, MCF = 1.40	65.59	34.48
	Consumption tax, MCF = 1.60	74.91	39.36
Baselines	Pure waste	0.00	0.00
	Welfare maximizing	100.00	100.00

Notes: Calculations based on national SUSENAS data. Social welfare is normalized so that 0% represents the welfare if the costs were incurred but no benefits received and so that 100% represents the welfare if the costs were incurred and the benefits were distributed in such a way as to maximize the social welfare, subject to the constraint that all individuals who received rice in each district received the same size transfer. "No program" represents the social welfare in the absence of the program, computed by multiplying the programs total cost by the marginal cost of funds shown, and allocating that cost across households proportionally to household consumption. Given these normalizations, the welfare level in the absence of the program increases as the program's welfare cost increases.

that the household received additional consumption with a value equal to the quantity of rice received multiplied by the value of the subsidy.[11] I assume that any missing rice went to wealthy corrupt officials, and so assign the consumption from all of the missing rice to the sampled household in each district with the highest per-capita expenditure. To calculate the social welfare in the absence of the program, I assume that no rice was allocated, but that the cost of the program, plus any dead-weight losses associated with revenue collection, was instead added back to household consumption in proportion to consumption.[12]

I normalize the social welfare so that 0 represents the social welfare had the taxation and costs for the program been incurred but no rice received, which I call the "pure waste" case, and so that 100% is the social welfare level that would have been achieved had the costs been incurred and the rice distributed in such a way as to maximize social welfare in each district, subject to the constraint that each household in a district receiving rice received the same amount, which I call the "welfare maximizing" case.

---

[11] The value of the subsidy is equal to the market price of rice, less the average price paid per kilogram of subsidized rice (Rp1050, or US$0.12) (LP3ES, 2000).

[12] Like many developing country governments, Indonesia raises the revenue for the OPK program through indirect taxes, whose incidence I assume is proportional to consumption. Estimates of the marginal cost of public funds for indirect taxes vary considerably, from a low estimate of 1.04 to 1.05 in Indonesia and Bangladesh (Devarajan et al., 2002) to between 1.17 and 1.56 in the United States (Ballard et al., 1985) to between 1.59 and 2.15 for India (Ahmad and Stern, 1987). For the administrative cost of implementing the program, I use the official government estimate of Rp137/kg. (Tabor and Sawit, 1999).

5.1. Results

The results are presented in Table 4. The social welfare actually achieved by the program is presented in row 1 of Table 4. The results suggest that, depending on the coefficient of relative risk aversion, between one-third and one-half of the potential welfare gain from the program was actually achieved. These welfare losses are attributable to two causes—corruption and imperfect targeting of benefits to poor households.

In row 2, I present the present the social welfare level that would have been achieved with the same set of beneficiaries, but with no corruption. This calculation allows us to separate the losses due to corruption from the losses due to imperfect targeting. The difference between rows one and two indicates the welfare cost of corruption in the program. The results suggest that corruption accounted for approximately 20% of the foregone welfare gain from the program, while imperfect targeting accounted for the remainder.

Of course, targeting is notoriously difficult, particularly in developing countries where data on incomes are unavailable, so achieving the socially optimum level of targeting may not be possible. In row 3, I calculated the welfare allocation under the official program eligibility guidelines, which was a simple proxy targeting scheme based on easily observable criteria.[13] Even though the correlation between the list of those eligible and those who actually received the rice was only 0.2, suggesting substantial local deviations from the official guidelines, the results suggest that allowing local control created almost no welfare change. Thus, while targeting may be important, simple fixes such as allowing local discretion may not provide the answer.[14]

The remaining rows of Table 4 present the social welfare level under the counterfactual that there was no program. The results suggest that when $\rho=1$, welfare would have been greater without the program if the marginal cost of public funds was greater than 1.12. On the other hand, had there been no theft of funds, the program would have been welfare increasing if the marginal cost of public funds was less than 1.35 (i.e., 35 cents of dead-weight loss for every $1 in revenue raised). Using a value of $\rho=2$, the program would have been welfare-increasing, even with corruption, up to a marginal cost of public funds of 1.42. At higher levels, the program would be welfare-increasing without corruption, but welfare decreasing with corruption. Thus, for reasonable values of the marginal costs of public funds, the estimates imply that without corruption, the program would have been welfare increasing, but given actual corruption levels, was in fact welfare decreasing.

---

[13] A household was initially eligible if it failed to meet any of the following five criteria: has a non-dirt floor, each household member eats at least 2 meals a day, each household member has at least 3 changes of clothes, children receive are brought to the doctor when sick, and households are able to fulfill basic religious obligations. Official eligibility was subsequently expanded somewhat, but I use these five criteria in computing the allocation in row 4 of Table 3.

[14] This finding is consistent with the results of Galasso and Ravallion (2005), who found that, in the Bangladesh Food-for-Education program, local autonomy created only mildly pro-poor targeting.

6. Conclusion

This paper has used data from a large transfer program in Indonesia to investigate the extent of corruption, and to see how the costs of corruption compare with the potential benefits from redistribution. I find that corruption is substantial—the central estimate is that at least 18% of the subsidized rice in the Indonesian program I study went missing. Corruption appears to be concentrated—over 60% of the missing rice comes from just 10% of the villages. Missing rice was more likely to be found in places that were ethnically heterogeneous and had lower population densities.

The estimates suggest that corruption in developing countries such as Indonesia may substantially inhibit a government's ability to carry out redistributive programs, particularly in rural areas. In the case of the Indonesian program studied here, for reasonable parameterizations of a social welfare function and assumptions for the marginal cost of public funds, the amount of corruption was substantial enough to make a program that would have been welfare enhancing become welfare reducing on net.

These results have important implications for the ability of developing countries to redistribute among their citizens. In cross-country work, LaPorta et al. (1999) find countries with higher perceived levels of bureaucratic corruption appear to have smaller transfer programs as a percent of GDP. The analysis in this paper provides complementary micro-level evidence that the costs of corruption can, in fact, outweigh the potential welfare benefits from redistribution.

Appendix A. Estimating the percentage of households receiving rice each month

I use the panel aspect of the SSD data to estimate the degree to which receipt of rice varied for individuals from month to month, and therefore, to correct for the fact that not all households received rice every month. Specifically, assume that in each month in each village, a proportion p_j of households receive rice, where j represents the village. This is the proportion that we would like to observe in order to calculate the total amount of rice received by households in each village each month. However, rather than observing p_j, we observe the percentage of households reporting receiving rice at any time during the previous three months, which I denote by r_{3j}. To recover p_j from r_{3j}, we need to know the

percentage of households on the recipient list who are replaced each month. Denote this percentage by α, and assume that each household on the recipient list has an equal probability of being replaced each month.[15] For a three-month window of observation, this assumption implies the following relationship between p_j and r_{3j}:

$$r_{3j} = p_j + \alpha p_j + \alpha p_j \left(\frac{1 - p_j - \alpha p_j}{1 - p_j} \right) \tag{2}$$

This expression simply states that the percentage of households reporting receiving rice at any point in a three month period is the percentage receiving rice in the first month, plus the percentage receiving rice in any subsequent month that had never received rice before. Note that the term in parentheses in Eq. (2) is to avoid double-counting the households that received rice in months one and three, but not in month two-when α and p_j are small, this second-order term will become negligible. For a four-month window of observation, the equivalent of expression (2) is:

$$r_{4j} = p_j + \alpha p_j + \alpha p_j \left(\frac{1 - p_j - \alpha p_j}{1 - p_j} \right) + \alpha p_j \left(\frac{1 - p_j - \alpha p_j - \alpha p_j \left(\frac{1 - p_j - \alpha p_j}{1 - p_j} \right)}{1 - p_j} \right) \tag{3}$$

For a given level of α and r_{3j} or r_{4j}, Eqs. (2) and (3) can be solved algebraically to yield the corresponding level of p_j.

It is possible to estimate the empirical value of α, the probability that a household receiving rice in one month did not receive it in the subsequent month, by using the panel aspect of the data from the SSD. In particular, we can match the actual correlation of rice receipt by particular households across waves of the survey with the correlation that would be implied by different levels of α. In fact, across the entire sample, the actual correlation between a household's reporting receipt of OPK rice in the May 1999 and October 1999 waves of the SSD is 0.44. This implied cross-wave correlation corresponds with a value of α of approximately 0.15—i.e., on average each month, 15% of households that received rice in the previous month were replaced by new households. This estimate of α can then be used to estimate the amount of rice actually received by households in each village in each month.

Appendix B. Statistical tests on the distribution of missing rice

I consider two statistical tests for the hypothesis that the missing rice is driven by under-reporting. To be conservative, in both tests, I ignore the fact that there is actually *over-*

---

[15] Given that poor households are more likely to receive the rice than wealthier households, the transition probability α is likely to depend on household characteristics, such as household income per capita. I abstract from this effect here, though I conjecture that including these effects would not significantly affect the estimated amount of corruption in the village.

reporting of rice in most villages, as discussed above. Instead, for the first test, I assume that each household in a village has an equal probability of receiving the rice in that village, and that there is constant underreporting across villages. Specifically, let r_j be the total number of kilograms of rice to be distributed in village j, N_j be the number of households in village j, and m be the overall estimate of the percent of rice missing—i.e., 17.8%. Then I assume that the probability a given household i receives rice is equal to $\frac{r_j}{20N_j}(1-m)$. I perform Monte Carlo simulations using this probability, and generate the joint distribution of both the mean amount of rice missing (μ) and the gini coefficient of the missing rice across villages (g) — i.e., the degree to which it is concentrated. Define the CDF of this joint distribution, which we estimate via the Monte Carlo simulations, as $F(\mu, g)$. The p-value of the statistical test is equal to $1 - F(\mu, g)$ at the actual levels of μ and g observed in the data.

The second test allows the probability that a household reports rice to depend on observable characteristics. The potential concern is that certain types of households might be more likely to underreport than others, and that these households might be concentrated in certain types of villages. I first perform a Probit regression of the probability a household i reports receiving the rice, as a function of the following variables: the household's per-capita consumption, size, number of children, age of household head, whether the household was headed by a woman or a widow, whether the household was eligible for the rice, whether the household owned a radio, TV, satellite dish, refrigerator, motorbike, car, dummies for type of roof, wall, flooring, toilet, and whether the household had electricity, and whether the household head was a widow, sick, illiterate, recently laid off, or illiterate. Denote the predicted probability from this regression as p_i, and the mean as p. I then assume that the probability a given household i receives rice is equal to $\frac{p_i}{p}\frac{r_j}{20N_j}(1-m)$, and proceed as above. This has the same overall probability of receiving the rice as the first test, but adjusts for the fact that certain types of households may be less likely to report receiving the rice, and that these households may be more or less clustered in particular villages.

References

Ahmad, E., Stern, N., 1987. Alternative sources of government revenue: illustrations from India, 1979–80. In: Newbery, D., Stern, N. (Eds.), The Theory of Taxation for Developing Countries. Oxford University Press, New York.

Ballard, C.L., Shoven, J.B., Whalley, J., 1985. General equilibrium computations of the marginal welfare cost of taxes in the United States. American Economic Review 75, 128–138.

Devarajan, S., Thierfelder, K.E., Suthiwart-Narueput, S., 2002. The marginal cost of public funds in developing countries. In: Fossati, A., Wiegard, W. (Eds.), Policy Evaluations with Computable General Equilibrium Models. Routledge Press, New York.

Di Tella, R., Schargrodsky, E., 2003. The role of wages and auditing during a crackdown on corruption in the city of Buenos Aires. Journal of Law and Economics 46, 269–292.

Fisman, R., Wei, S., 2004. Tax rates and tax evasion: evidence from 'Missing Imports' in China. Journal of Political Economy 112, 471–496.

Galasso, E., Ravallion, M., 2005. Local knowledge vs. local accountability? Decentralized targeting of an anti-poverty program. Journal of Public Economics 89, 705–727.

Glaeser, E.L., Saks, R.E., 2004. Corruption in America. NBER Working Paper #10821.

LaPorta, R., Lopez-de-Silanes, F., Shleifer, A., Vishny, R., 1999. The quality of government. Journal of Law, Economics, and Organization 15, 222–279.

LP3ES, 2000. Laporan Akhir Studi Evaluasi Operasi Pasar Khusus Beras (Final Report: Evaluation Study of OPK Rice Program). Mimeo. Jakarta.

Mauro, P., 1995. Corruption and growth. Quarterly Journal of Economics 110, 681–712.

Olken, B.A., 2005. Monitoring corruption: evidence from a field experiment in Indonesia. Mimeo. National Bureau of Economic Research.

Olken, B.A., Nabiu, M., Toyamah, N., Perwira, D., 2001. Sharing the wealth: how villages decide to distribute OPK rice. SMERU Working Paper, Jakarta.

Pritchett, L., Sumarto, S., Suryahadi, A., 2002. Targeted programs in an economic crisis: empirical findings from the experience of Indonesia. Mimeo. Harvard University.

Reinikka, R., Svensson, J., 2004. Local capture: evidence from a central government transfer program in Uganda. Quarterly Journal of Economics 119, 678–704.

Rose-Ackerman, S., 2004. The challenge of poor governance and corruption. Mimeo. Copenhagen Consensus and Yale University.

Tabor, S.R., Sawit, H.M., 1999. The OPK program: economy wide impacts. Mimeo. Indonesian Ministry of Food.

Treisman, D., 2000. The causes of corruption: a cross-national study. Journal of Public Economics 76, 399–457.

Yang, D., 2003. How easily do lawbreakers adapt to increased enforcement? Philippine smugglers' responses to a common customs reform. Mimeo. University of Michigan.