Corruption and the Costs of Redistribution*

Micro Evidence from Indonesia

Benjamin A. Olken

Harvard University

July 2003

ABSTRACT

This paper examines the degree to which the corruption in developing countries may impair the ability of governments to redistribute 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 of program rice actually received by households. The central estimates suggest that, on average, 18 percent of the rice appears to have disappeared. 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.

* I wish to thank Abhijit Banerjee, Francesco Caselli, David Cutler, Esther Duflo, Amy Finkelstein, Ed Glaeser, Caroline Hoxby, Larry Katz, Michael Kremer, Seema Jayachandran, Lant Pritchett, Andrei Shleifer, and Sudarno Sumarto for helpful comments. I also wish to thank the staff of SMERU, particularly Asep Suryahadi, Wenefrida Widyanti, and Sri Kusumastuti Rahayu, for invaluable assistance in Indonesia, and to acknowledge financial support from the Social Science Research Council Program in Applied Economics. All views expressed, however, are my own. Email: bolken@fas.harvard.edu.

I. Introduction

A major distinction between developed and developing countries is the endemic amount of corruption in much of the developing world. A number of authors have suggested that corruption may have substantial efficiency consequences, as it can reduce protection for property rights, inhibit the government's ability to solve market failures, and create substantial uncertainty in the business climate.[1] Indeed, it has been suggested that corruption may be a major contributor to the poor growth rates observed in many developing countries (Mauro 1995).

In this paper, I explore another potentially important – but largely overlooked – implication of corruption: its potential to inhibit redistribution. Transfer programs and social safety nets have the potential to substantially improve welfare in the developing world. (Ravallion 2003) 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.

I examine the costs of corruption empirically in the context of 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. As a result, empirical work on corruption has been dominated by the use of subjective assessments of corruption. In this paper, however, I build on a small but growing literature that seeks to measure corruption by comparing two measures of the same

[1] For an overview of the literature on corruption, see Bardhan (1997).

quantity, one “before” and one “after” the corruption takes place.[2] 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 comparing these records with household surveys on the amount of subsidized rice actually received by villagers in those villages.

Using this approach, I find that at least 18 percent of the program’s rice disappeared between the time it left government warehouses and the time it reached households. This estimate constitutes a lower bound, and as discussed in the text, the actual amount of corruption may be substantially higher. The missing rice appears to be highly concentrated—I detected missing rice in just over one third of villages, but in those villages missing rice, the mean amount of rice missing was 43%.

In the second part 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.

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 that perceptions of country’s level of bureaucratic corruption appear to be inversely related to the size of a country’s transfer programs. 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. Furthermore, these results suggest that by not accounting for corruption, recent empirical work evaluating the effects of decentralizing redistribution programs in developing countries, such as Alderman (2001) and Galasso and Ravallion

[2] For other examples of this approach, see Reinikka and Svensson (2002), who examine corruption in educational expenditures, and Fisman and Wei (2001) and Yang (2003), who examine corruption in international trade.

(2002), may be missing a determinant of program success of similar magnitude to the differences in targeting on which they focus.

The remainder of the paper is organized as follows. Section II discusses the background of the Indonesian poverty alleviation program I study, known as the *OPK* program. Section III presents the estimates of the extent of corruption in the program. Section IV presents the welfare calculation used to compare the costs of corruption with the benefits from redistribution. Section V concludes.

II. Background: Corruption and the *OPK* Program

In 1997-1998, Indonesia experienced a severe economic collapse. In one year, the value of the local currency fell by over 80% and real GDP fell by 13%. As a result, the percentage of people living below the poverty line increased from 15 percent before the financial crisis to 27 percent in 1999 (Pradhan et. al, 2000). 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 20kg of OPK rice per month, at a price approximately 60% below market.[3] The size of this subsidy was substantial—for the median eligible household purchasing their full allotment of rice, the subsidy represented approximately 10.2% of total pre-program monthly household expenditures.[4] Approximately 50 percent of rural Indonesian households were eligible to receive the subsidized OPK rice.

[3] As local rice markets are quite liquid, households who did not need the full 20kg of subsidized rice could still buy the subsidized rice and resell it at market prices to private traders, though some of the value of the subsidy would be lost to transaction costs. (See Olken et. al 2001 for an example of this from the field.)
[4] Author's calculations using the SUSENAS dataset.

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 20kg 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 one village I visited in August 2001, residents reported that the rice deliveries had become intermittent in 2000, and had stopped completely by 2001. Yet, the government warehouse reported that 1.6 tons of rice had been distributed to the village on time every month, including the month prior to my visit. This figure was confirmed by the subdistrict chief and by a village head from a neighboring village, but the village head from the village in question refused to discuss OPK at all. (See Olken et. al 2001 for more information.) Based on the available evidence, it seemed as though the village head, 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.

More common than this type of all-out theft, however, may have been other methods which may have enabled a village head to retain some fraction of the rice. 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. In December 1998, for example, the per-household allocation of rice was raised from 10kg/month to 20kg/month, though news of this change was not well publicized at the village level. It is easy to see how, by not passing on the change in program rules to villagers, a village head could end up with 50% of the program funds with relatively little fear of public notice. Perhaps more common, in many areas the allocation was reduced from 10 (or 20) *kilograms* per household to 10 (or 20) *liters* per household. Since there are approximately 0.8 kilograms of rice in a liter, by changing from kilograms to liters and keeping the number of recipient households fixed, a village head could generate a surplus equal to approximately 20% of the village's rice allocation. The "missing rice" I detect in the empirical work below likely includes rice obtained through all of these mechanisms, as well as many others.

III. Estimating the amount of "Missing Rice"

Detecting corruption is almost always difficult, as corrupt actors go to great lengths to cover their tracks. While it is not possible to detect who specifically stole OPK rice, it is possible to estimate the total amount of rice which appears to have disappeared from the program in a particular village. To do this, one needs to estimate two quantities: the total amount that was *disbursed* by the central government to each village and the total amount of rice that was *received* by villagers in each village. 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.

Estimating the amount of rice received by villagers

To estimate the amount of rice received by villagers, I use household level survey data from the 1998-1999 Hundred Villages Survey, known in Indonesian as the *Survei*

Seratus Desa, or SSD. The SSD is a four-wave panel dataset conducted in 100 poor communities throughout Indonesia by the Indonesian Central Statistics Office. Each wave of the SSD covered 120 households in each village, and is representative at the village level. Summary statistics from the SSD are presented in Table 1.

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.[5] 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 for 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.

There are, however, two issues with using this survey to estimate the amount of rice received by villagers. First, the SSD only provides information on *whether* the individual receives rice. It does not provide 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 20kg monthly allotment. Quantitative and qualitative evidence suggests that this assumption is generous, and therefore likely to produce an underestimate of the amount of corruption. In particular, a 1999 Indonesian survey found that only 19% of households received the full 20kg, 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 households received substantially less than 20kg of rice, and never found evidence of any households receiving

[5] 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.

[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.

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, in one wave, four months, 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.

To correct for this, I assume that a fixed percentage of households on each village's OPK recipient list are replaced each month. I denote this percentage by α, and assume that each household on the recipient list has an equal probability of being replaced each month. I then use the panel aspect of the data to estimate the value of α that best corresponds to the data. I discuss the details of this correction, and the estimation of α using the panel aspect of the SSD, in Appendix 1. I find that a value of α of 0.15 appears to best fit the data—i.e., it appears that each household that received rice in a given month had an 85% chance of receiving it in the subsequent month. In the discussion of the results, I therefore focus on results when α is equal to 0.15, though I also report results for alternative levels of α. In particular, I report the results where $\alpha = 0$, which corresponds to the most conservative assumption that a household reporting receipt of rice in *any* of the previous three months received it in *every one* of the three months.

Estimating the amount of rice disbursed to each village

For information on the amount of rice that each village *should have* received, I match the villages in the SSD with administrative records on the amount of rice disbursed to the village. While information on the total amount distributed to each district warehouse by the central government was available from the central government in Jakarta (see Tabor & Sawitt

1999), information on the amount distributed to individual villages was kept only at the district level distribution centers (known as DOLOGs and Sub-DOLOGs). Accordingly, for each of the 100 villages in the SSD, I obtained from each district level distribution office the number of kilograms of rice distributed each month to each village in the SSD.

As discussed above, the household data in the SSD on rice receipt asked households if they received any rice during the three or four months prior to the survey. Accordingly, for each of round of the SSD survey, I compare the average amount of rice distributed in the village in any of the three or four months prior to the survey to the estimate computed above for the maximum amount of rice received by households. Only when the maximum amount of rice received by households is less than the amount disbursed to the village do I conclude that rice was missing from that village.

Results

In Table 2, I present the results from this procedure. The columns of Table 2 represent different assumptions for the maximum amount of rice a household could receive; the rows of Table 2 represent different assumptions for α, the probability that a household receiving rice in one month receives rice in the subsequent month. The column and rows in bold indicate the assumptions I focus on, i.e., the assumption that no household received more than 20kg of rice each month that that α, the proportion of households that received rice in one month but not in the subsequent month, was equal to 0.15.

Under these assumptions, I estimate that 17.8 percent of the rice appears to be missing. Even assuming that α is equal to 0—the most conservative case discussed above—11.1 percent of the rice would still be missing. As discussed above, the 17.8 percent estimate may still be substantially below the actual amount of missing rice. If, for example, the maximum amount households were assumed to receive was 15 kg instead of 20 kg per month, the estimate of the percentage rice missing jumps to 28.0 percent.

An interesting question is whether all of this missing rice is coming from just a few villages in which all of the rice is missing, or whether the theft was spread more evenly across many villages. Table 3 presents estimates of the percent of village-wave observations in which I detect missing rice. Under the baseline assumptions, the data indicate that there was missing rice in 37 percent of villages. Even in the most conservative case with $\alpha = 0$, fully 21 percent of villages were estimated to have some missing rice.

This suggests that the missing rice I detect is relatively concentrated—just over a third of all villages were missing rice. Furthermore, even within those villages missing rice, the majority of the missing rice comes from a small subset of villages missing substantial amounts. Figure 1 presents the distribution of the missing rice among the 37 percent of villages where at least some rice was missing. Figure 1 shows, for example, that 65 percent of the villages missing rice were missing at least 25% or more of their total allocation, and that the mean amount of rice missing, conditional on missing any rice, was 43% of the allocation.

A potential concern is that the estimates of missing rice are driven by underreporting of OPK receipt by households in the survey. The concentration of the missing rice, however, suggests that this is not the case, unless for some reason underreporting was very non-uniform across few villages. To see this, note that if the chance 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.

Since the villages in the SSD are poorer than the typical Indonesian village, one needs to proceed with some caution when extrapolating from these results to the entire country. In

Appendix 2, I discuss how I use the nationally-representative SUSENAS, matched with district-level administrative data on OPK obtained from Tabor and Sawitt (1999), to perform a similar calculation at the district level. The results from the SUSENAS are consistent with the results in the SSD, and in fact suggest that corruption in the program nationally may have been slightly higher than in the 100 villages sampled in the SSD.

An natural next step would be to investigate the correlates of corruption—i.e., which social, economic, and political factors appear to be related to higher levels of corruption. Unfortunately, in this context data limitations make doing so difficult.[7] I regard this as a natural direction for future work in a different setting.

IV. 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 percent 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.

[7] 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 substantially increase the number of households receiving rice. Using the amount of missing rice as a dependent variable would therefore conflate factors related to theft with factors related to concentrating rice among a small number of recipients. With the existing data, separating these two effects would be impossible without implausible exclusion restrictions.

Methodology

Welfare analysis rests on an assumption of a social welfare function. I focus on a utilitarian social welfare function with CRRA individual utility, i.e. a function of the following form:

$$W = \frac{1}{N}\sum_{i=1}^{N}\frac{y^{1-\rho}}{1-\rho} \tag{1}$$

where y represents household expenditure per capita and N represents the number of households. I present results for the case where $\rho = 1$ (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 three different social welfare levels: the social welfare level actually achieved by the program, the social welfare that would have achieved by the program in the absence of corruption, and the social welfare that would have been attained in the absence of the program. These social welfare estimates are based on the actual joint distribution of household expenditure and rice receipt in the data.

To calculate the social welfare level actually achieved by the program, for each household in the data that reported receiving OPK rice, I assume that the household received additional consumption with a value equal to the quantity of rice received multiplied by the value of the subsidy.[8] I assume that any missing rice went to the wealthy, as corrupt officials, such as village heads, are among the wealthiest in an area. I therefore assign the consumption from all of the missing rice to the sampled household in each district with the highest per-capita expenditure.

[8] Note that this assumes no substitution effects. This seems reasonable since, for virtually all households in the sample, the total rice consumption even without OPK rice was above the total amount they were eligible to buy from the OPK program. The program therefore did not affect the marginal price of rice faced by most households.

To calculate the social welfare that would have been achieved by the program in the absence of corruption, I perform the same calculation, but with a different treatment of the missing rice. Rather than assume that all of the missing rice ended up in the hands of a corrupt official, I calculate the social welfare assuming that the missing rice was divided up equally among the households that actually received program rice. The difference between this allocation and the previous allocation allows us to calculate the welfare cost of corruption, holding the targeting of the program fixed.

Finally, 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 the incidence of the tax used to pay for the program. The deadweight losses and tax incidence of the program costs depends how the Indonesian government raises its revenue. Like many developing country governments, Indonesia raises the revenue for the OPK program through indirect taxes, whose incidence I assume is proportional to consumption. I therefore assume that each household's share of the total program cost is proportional to that household's expenditure.

In addition to these three cases, I calculate two other social welfare levels for use as benchmarks. First, I calculate the social welfare that would have been achieved had the taxation and costs for the program been incurred but no rice received, which I call the "pure waste" case. Second, I calculate the social welfare 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 village, subject to the constraint that each household receiving rice received the same amount, which I call the "welfare maximizing" case. I normalize the pure waste case to 0% and the welfare maximizing case to 100%, so that the magnitudes of other social welfare levels are more easily interpretable.

Implementation

Implementing this procedure requires several parameters: the value of the subsidy, the marginal cost of public funds (used for calculating the dead-weight loss from taxation), the administrative costs of running the program, and the joint distribution of household expenditure and OPK receipt in Indonesia. This section describes the estimates I use for each of these parameters.

To obtain the value of the subsidy, I subtract the average price paid for the subsidized rice from the average market price of rice. According to the 1999 survey of OPK recipients, the average price paid per kilogram of subsidized rice was Rp. 1,050 (US$0.12), which represents a premium of Rp. 50/kg over the official subsidized price. (LP3ES, 2000) This markup may reflect the additional transport costs incurred by villages in distributing the rice that were passed on to households. The average market price for rice is each district is reported in the data. The difference between these two quantities represents my estimate for the value of the subsidy.

Estimates of the marginal cost of public funds for indirect taxes vary considerably depending on the country and the method used. For developing countries, estimates for the marginal cost of funds from indirect taxes range from a low estimate of 1.04 to 1.05 in Indonesia and Bangladesh (Devarajan et. al, 1999) to between 1.59 and 2.15 for India (Ahmad and Stern, 1987). By comparison, estimates of the marginal cost of public funds for the U.S. range from 1.17 to 1.56 (Ballard et. al, 1985), with policy analysis typically using values in the range from 1.30 to 1.40. I therefore report results for a wide range of assumptions about marginal costs of public funds, including 1.00 (i.e., no dead weight loss). For the administrative cost of implementing the program, I use the official government estimate of Rp. 137/kg. (Tabor and Sawitt 1999)

To ensure that the sample I use is for the welfare calculation is representative of Indonesia, I perform the calculation using data on household expenditure and rice receipt from the nationally-representative 1999 SUSENAS survey. The advantage of using the SUSENAS, rather than the 100 villages in the SSD, is that the national sample in the SUSENAS allows me to account for net redistribution from rich urban areas to poorer rural areas. As discussed in Appendix 2, however, direct estimates of corruption using the SUSENAS are likely to miss a substantial amount of corruption that would be detectable using the village-level approach I use in Section III. To adjust for this under-reporting, I scale the total amount of estimated corruption in each district in the SUSENAS up so that the total amount of corruption nationwide is equal to the 17.8 percent I estimated using the SSD.

Results

The results are presented in Table 4, where the social welfare from the benchmark “pure waste” case is normalized to 0 and the benchmark social welfare from the “welfare maximizing” case is normalized to 100%. 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, whereas the difference between row two and the welfare-maximizing level (100%) indicate the welfare cost of targeting failures. The results suggest that corruption accounted for approximately 20% of the foregone welfare gain from the program, while imperfect targeting accounted for the remainder.

The remaining rows of Table 4 present the social welfare level under the counterfactual where 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). Thus, for reasonable values of the marginal costs of public funds (between 1.12 and 1.35), the estimates imply that corruption resulted in a program that would have been welfare increasing becoming welfare decreasing.

These results change somewhat as the coefficient of risk aversion increases. 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 becomes welfare-increasing without corruption, but welfare decreasing with corruption. Though this threshold cost of public funds is higher than the point-estimate for Indonesia from Devarajan et. al (1999), it is well within the range estimated for other developing countries. At higher levels of the coefficient of relative risk-aversion, the threshold marginal cost of public funds at which corruption would make the program welfare-decreasing would also be correspondingly higher.

Clearly, the specific conclusions presented in this section about whether the program was welfare increasing or decreasing depend substantially on the particular assumptions made about the marginal cost of public funds, the coefficient of relative risk aversion, the functional form of the utility function, and the social welfare function. The results should therefore be considered suggestive. Nevertheless, this section has demonstrated that that corruption on the order of 18% is substantial, large enough so that, under plausible parameterizations and assumptions, it can turn a beneficial program into a program that is actually welfare-decreasing. It has also shown that the losses from corruption are large on an absolute scale, and that an evaluation of redistributive programs that fails to include the costs of corruption may miss a quantitatively important factor in determining program success.

V. 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 percent of the subsidized rice in the Indonesian program I study went missing. Corruption appears to be concentrated—just over a third of villages were missing any rice, but those villages that were missing rice were missing an average of 43 percent.

The estimates suggest that corruption in developing countries such as Indonesia may substantially inhibit a government's ability to carry out redistributive programs. 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 findings have important implications for thinking about corruption more generally. A substantial literature has demonstrated that corruption may negatively impact efficiency. The analysis here suggests that future work on corruption should also consider the effects of corruption on redistribution when evaluating the welfare impact of institutional quality. In addition, the costs of corruption for implementing redistribution programs may also help explain the cross-country correlation between the level of transfers in a country and the external perceptions of a country's institutional quality, and more generally, the low levels of transfer programs in developing countries.

References

Ahmad, Etisham and Nicholas Stern, "Alternative Sources of Government Revenue: Illustrations from India, 1979-80," in David Newbery and Nicholas Stern eds., *The Theory of Taxation for Developing Countries*, New York: Oxford University Press, 1987.

Alderman, Harold, "Do local officials know something we don't? Decentralization of targeted transfers in Albania," *Journal of Public Economics 83 (3)*, pp. 375-404, March 2002.

Ballard, Charles L., John B. Shoven, and John Whalley, "General Equilibrium Computations of the Marginal Welfare Cost of Taxes in the United States," *American Economic Review 75 (1)*, pp 128-138, March 1985.

Bardhan, Pranab, "Corruption and Development: A Review of Issues," *Journal of Economic Literature 35 (3)*, pp 1320-46, September 1997.

Devarajan, Shantayanan, Karen E. Thierfelder, and Sethaput Suthiwart-Narueput, "The Marginal Cost of Public Funds in Developing Countries," in Amedeo Fossati and Wolfgang Wiegard (eds.), *Policy Evaluations with Computable General Equilibrium Models*, New York: Routledge Press, 2002.

Fisman and Wei, "Tax Rates and Tax Evasion: Evidence from 'Missing Imports' in China," NBER Working Paper #8551, 2001.

Galasso, Emanuela, and Martin Ravallion, "Local Knowledge vs. Local Accountability? Decentralized Targeting of an Anti-Poverty Program," *Journal of Public Economics*, forthcoming.

LaPorta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert Vishny, "The Quality of Government," *Journal of Law, Economics, and Organization 15 (1)*, pp. 222 – 279, 1999.

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

Mauro, Paulo, "Corruption and Growth," *Quarterly Journal of Economics*, 1995.

Olken, Benjamin A., Musriyadi Nabiu, Nina Toyamah, and Daniel Perwira, "Sharing the Wealth: How villages decide to distribute OPK rice," SMERU Working Paper, Jakarta, 2001.

Pradhan, Menno, Asep Suryahadi, Sudarno Sumarto, and Lant Pritchett, "Measurements of Poverty in Indonesia: 1996, 1999, and Beyond,"

Pritchett, Lant, Sudarno Sumarto, and Asep Suryahadi, "Targeted Programs in an Economic Crisis: Empirical Findings from the Experience of Indonesia," mimeo, Harvard University, 2002.

Ravallion, Martin, "Targeted Transfers in Poor Countries: Revisiting the Tradeoffs and Policy Options," World Bank Working Paper #3048, 2003.

Reinikka, Ritva and Jakob Svensson, “Explaining Leakage of Public Funds,” World Bank Working Paper Number 2709, 2002.

Tabor, Steven R., and Sawit, Husein M., “The OPK Program: Economy Wide Impacts,” Indonesian Ministry of Food, mimeo, 1999.

Yang, Dean, “How Easily do Lawbreakers Adapt to Increased Enforcement? Philippine Smugglers' Responses to a Common Customs Reform,” Mimeo: University of Michigan, 2003.

Appendix 1: Estimating the percentage of households receiving rice each month

I use the panel aspect of the 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.[9] 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 equation (2) is to avoid double-counting the households that 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}, equations (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 percent 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.

[9] 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.

Appendix 2: Comparing corruption in the 100 villages to corruption in the rest of Indonesia

To estimate the degree to which the estimate of corruption from the SSD is representative of corruption in Indonesia as a whole, I perform a calculation similar to that in Section III using data from the nationally-representative 1999 SUSENAS, the national social welfare survey. The SUSENAS includes data on the number of times each household received OPK rice during the six months prior to the survey, but like the SSD, collects no information on the quantity of rice received by each household.[10]

Unlike the SSD, which contains a representative sample of each village, the SUSENAS is only representative at the district level.[11] In principle, if the data contained information on the quantity of rice received by each recipient, the procedure I use to determine the amount of missing rice would be equally valid using a district level sample as a village level sample. However, because quantity data is unavailable, I again make the conservative assumption that each recipient received the maximum amount of rice possible. Making this assumption over an entire district, rather than a single village, will tend to substantially under-estimate the amount of missing rice.

To see why this is the case, note that while some village heads stole rice, other village heads divided the rice into many small portions, which they then distributed widely. In some villages, for example, every household reported receiving some rice, even though the amount they received was only a few kilograms per household. (See, e.g., Olken et. al. 2001) By aggregating over multiple villages in a district, and assuming that each households that received rice received the full 20 kilograms, the villages where many households received a small amount of rice will mask the theft of the rice in other villages in the same district. For this reason, the estimates of missing rice done on a village-by-village basis are more accurate than those at the district level, which is why I focus on the results from the SSD rather than the SUSENAS.

To compare the results from the nationally-representative SUSENAS to those from the SSD, then, it is necessary to re-do the calculations in the SSD aggregating the data to the district level, rather than the village level. This estimate of missing rice will be substantially smaller than the estimate of missing rice obtained by looking at the village level, but will be comparable to the number estimated using district-level data from the SUSENAS. One can then compare these two calculations to see whether, at the same level of aggregation, there appears to be more or less missing rice in the SSD areas than in the overall national sample.

In fact, the estimates of missing rice obtained by aggregating at the district level are remarkably similar in both the SSD and the SUSENAS, suggesting that SSD villages are not substantially different from the national average in their propensity to have missing rice. Aggregating the SSD data to the district level, I estimate that 8.3 percent of rice was missing, as opposed to the 17.8 percent estimate I calculated when the data was aggregated only to the village level. This suggests that aggregating to the district level rather than the village level misses 53% of the actual missing rice. Applying this ratio to the data from the SUSENAS suggests that a total of 19.8 percent of the rice was missing nationally—slightly higher than the estimate obtained from the SSD.

[10] Note that I drop the provinces of East Timor and Irian Jaya, as well as major cities (*kotamadya*), where the program was substantially different. The results from the SUSENAS are therefore representative of all of rural and semi-urban Indonesia, with the exception of the aforementioned provinces.

[11] At the time of the survey, Indonesia contained approximately 300 districts, with an average of 215 villages per district.

Table 1: Summary statistics from the SSD

Household Characteristics	
Log HH expenditure per capita	11.34
	(0.47)
HH size	4.17
	(1.78)
Does HH have electricity	0.591
	(0.492)
Did HH receive OPK rice?	0.561
	(0.496)
Was HH eligible to receive OPK rice?	0.668
	(0.471)
Village Characteristics	
Village population	3,655
	(2,568)
Percent population poor (KPS or KS1)	0.588
	(0.263)
Distance to nearest major town (km)	77.12
	(74.03)

Notes: Standard deviations in parentheses.

Table 2: Estimates of missing rice

		Estimated fraction of rice missing		
		Maximum amount of rice received by a household		
		15	**20**	25
	0.00	0.184	0.111	0.079
	0.05	0.218	0.132	0.091
Level	0.10	0.252	0.155	0.105
of α	**0.15**	0.280	**0.178**	0.120
	0.20	0.307	0.199	0.135
	0.25	0.328	0.219	0.152

Notes: Each cell represents the estimated share of rice unaccounted for using the assumptions listed the column and row, based on calculations using the SSD. The assumption in the column represents the maximum number of kg each household is assumed to have received in any given month. The assumption in the row (α) represents the fraction of households whose rice receipt is assumed to have alternated each month. This assumption is used to account for the fact that the question on rice receipt in the SSD asked about receipt during the previous 3 or 4 months, whereas the rice arrived each month. The rows and columns in bold (maximum rice = 20kg and α = 0.15) indicate the assumptions that appear to best fit the data.

Table 3: Percent of villages with missing rice

		Estimated fraction of villages with rice missing		
		Maximum amount of rice received by a household		
		15	**20**	25
	0.00	0.380	0.246	0.187
	0.05	0.433	0.292	0.205
Level	0.10	0.471	0.355	0.238
of α	**0.15**	0.512	**0.372**	0.267
	0.20	0.535	0.401	0.296
	0.25	0.567	0.427	0.345

Notes: See Notes to Table 2. Each cell represents the estimated percent of villages with at least some rice missing using the assumptions listed the column and row, based on calculations using the SSD.

Table 4: Comparing costs and benefits

Allocations:		*Utilitarian, CRRA utility* $\rho = 1$	*Utilitarian, CRRA utility* $\rho = 2$
Program	Actual allocation	52.23%	35.31%
	Actual allocation, no corruption	62.06%	42.73%
No Program	Consumption tax, MCF = 1.00	46.90%	24.68%
	Consumption tax, MCF = 1.05	49.24%	25.91%
	Consumption tax, MCF = 1.10	51.58%	27.14%
	Consumption tax, MCF = 1.20	56.25%	29.59%
	Consumption tax, MCF = 1.30	60.92%	32.03%
	Consumption tax, MCF = 1.40	65.59%	34.48%
	Consumption tax, MCF = 1.50	70.25%	36.92%
	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 in each village 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 given, 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.

Figure 1: Distribution of missing rice