

3.2. Does Corruption Matter?

Although the previous section has shown that corruption is substantial in magnitude—whether in the form of bribes given to civil servants, graft from public expenditures, or public employees not showing up to work—this does not necessarily answer the question of whether corruption actually has a negative impact on economic activity. Instead, one must test directly for whether corruption has efficiency costs.

For example, Gorodnichenko and Peter (2007) showed that, on average, public employees in Ukraine have the same consumption levels as their private sector counterparts, even though their salaries are 24-32 percent lower. Corruption in this case does not seem to be providing extra income to these public employees, as what the government pays them is reduced exactly to offset the amount they receive in bribes. Whether corruption “matters” for economic activity in this case does not depend on the amount of corruption: instead, the economic efficiency losses (or gains) from corruption depend on whether the deadweight loss imposed by the bribes they collect is greater than (or smaller than) the equivalent deadweight loss from taxation that would be needed to raise the revenue to pay the equivalent amount of money in salaries were corruption was eliminated.

More generally, corruption could have either efficiency costs or lead to efficiency gains. This section lays out the evidence thus far on three ways in which corruption may have aggregate efficiency costs: macro impacts, the costs imposed on firms, the costs imposed on government activity, and the costs imposed through the government’s lack of ability to correct externalities.

3.2.1. *Macro Impacts*

The macro literature on corruption has tried to estimate the effect of corruption on macro variables like growth, investment and foreign aid, although the empirical strategy used in these types of studies is usually not very well identified. Most of these studies are cross-country comparisons that rely on international survey and perception-based indices to provide some evidence on the negative effect of corruption at the aggregate level. For example, in an early and very well-known study Mauro (1995) examined the effect of corruption on economic growth. He shows a negative correlation between corruption and investment and growth.

Also using a cross-country comparison and firm-level dataset, Javorcik and Wei (2009) studied the impact of corruption on foreign direct investment in emerging markets. Specifically, the study distinguishes two effects of corruption: a reduction in the volume of foreign direct investment and a shift on firm ownership structure. The results show corruption reduces inward foreign direct investment and shifts the ownership structure towards joint ventures, probably as a response to minimize the transaction cost of dealing with corrupt governments. Similarly, D’Souza (2012) finds in the aftermath of the 1997 OECD Anti-Bribery Convention, bilateral exports from signatory countries declined by 5.7 percent to more corrupt importers. The convention indirectly increased the transaction costs of exporting to corrupt countries (who expect bribes) by imposing a credible punishment for bribing foreign officials. Although

foreigners could also reduce foreign aid as an instrument to correct misgovernment, Alesina and Weder (2002) provide evidence that corruption does not affect the allocation of foreign aid or debt relief.

While these studies—particularly the Mauro study—launched empirical work on the impacts of corruption, the challenge is that corruption is potentially endogenous, and finding credible instruments for corruption at the macro level is difficult. We therefore turn to micro evidence for more reliable estimates of the impact of corruption.

3.2.2. Impact on Firms

To estimate the efficiency cost of corruption on firm behavior, ideally one must know several things. First, one would like to estimate how corruption changes the effective marginal tax rate faced by firms. To the extent that bribery is used to reduce tax liabilities (e.g., bribing tax officials to reduce tax payments), one would expect the marginal bribe rate to probably be lower than the official marginal tax rate, so corruption reduces effective tax rates. On the other hand, if bribes are charged for other types of government activities, this could add to the effective marginal tax rate faced by firms.

Svensson's (2003) study of firm bribe-paying behavior in Uganda provides some clues that while there is a positive relationship between bribes and firm profits, it is very flat. Specifically, he estimates that each USD 1.00 in firm profits per employee leads to about USD 0.004 in additional bribes paid, for a "marginal bribe rate" of 0.4 percent on profits. He also finds that each USD 1.00 in capital stock per employee leads to an additional USD 0.004 in additional bribes paid, representing an additional 0.4 percent "marginal bribe rate" on capital stock. Note that these are marginal rates: the average level of bribes is substantially higher, but bribes increase relatively weakly with profits and capital stock. If the only impact of corruption was to impose a tax of 0.4 percent on profits and 0.4 percent on capital, one might expect that the impact of corruption on firm activity would be quite modest. By way of comparison, the marginal tax rate on corporate profits for large corporations in the United States is 35 percent.

Note that the Svensson study does not answer the question of the impact of corruption on firms. Instead, it only establishes what their effective corruption tax rates are. Furthermore, there may be other factors ways in which corruption could affect firm behavior beyond the marginal tax rate. For example, many have argued that the uncertainty surrounding corruption makes it more costly than an equivalently-sized tax. Wei (2000), for example, makes this argument looking at foreign direct investment and measuring uncertainty through perceptions-based metrics. More recently, Malesky and Samphantharak (2008) examine the impact of changes in governors in Cambodia. Using survey data, they show that these changes are associated with increases in uncertainty about corruption, but reductions in actual corruption levels and decreased firm-level investment.

Sequeira and Djankov (2009) examine a different type of distortion: changes in the firm's production choices designed to avoid corruption. Specifically, they estimate the economic costs and distortions associated to corruption acts at two ports in Mozambique and South Africa by directly observing bribe payments to port and border post officials for a random sample of 1,300 shipments. The study calculates the difference in the amount of bribes paid between the two ports and between goods that experienced a

tariff reduction and those that did not. The results show bribes are high, frequent and vary significantly across ports. Bribes can represent up to a 14 percent increase in total shipping costs for a standard container while accounting for a 600 percent increase in the monthly salary of a port official. Further, the study shows bribes do not only increase shipping cost but also impose high distortions on firms. Estimations suggest that even when accounting for distance and the urgency of the shipment among other variables, corruption is the strongest predictor of the choice of port. About 46 percent of South African firms located in regions in which overland costs to the port of Maputo are 57 percent lower go the long way around to Durban to avoid higher bribe payments. This represents a real efficiency loss: firms are willing to pay higher (real) trucking costs to avoid having to pay bribes in Mozambique.

Given that corruption could have both direct effects (through changing the effective marginal tax rate) as well as indirect effects (through uncertainty or other channels), it is necessary to examine the net impact of corruption on firm decisions directly. Fisman and Svensson (2007) perform this exercise in Uganda, using the same dataset as in Svensson (2003). They calculate both bribes and tax payments as a function of total firm sales, with bribes averaging 1.3 percent of sales and taxes averaging 8.5 percent of sales. They regress firm growth over a two year period on the bribe and tax rate, instrumenting for the bribe and tax rate with industry-by-location averages. They find that a 1 percentage point increase in bribes reduces annual firm growth by three percentage points. By comparison, a 1 percentage point increase in taxes reduces annual firm growth by 1 percentage point, so bribes have three times the negative impact of taxes on firm performance. They interpret the findings as showing that the negative impacts of bribes on firm activity are higher than the corresponding impacts of taxation—with substantially large magnitudes for both.¹

While this study provides one data point on the relationship between bribes and firm growth, substantial work remains to be done. In these studies, bribe rates are not randomly assigned, and there is also no clear estimation of the marginal bribe rate and the impact of bribes on marginal tax rates. Ideally, one would have plausibly exogenous variation in corruption levels across jurisdictions, which could then be used to estimate the impact of different types of corruption on firm behavior. Finding such a change and linking it to firm behavior, seems an important next step in understanding the efficiency costs of corruption.

Corruption could also imply that firms lack even the opportunity to compete on contracts to provide public services because government officials set up their own companies to enjoy the lucrative contracts. Understanding how this dynamic results in overpriced contracts and the under-provision of products and services is also an important part of the future research agenda.

3.2.3. Impact on Government Provision of Goods and Services

¹ Although the level of this effect seems enormous, it is worth recalling that the bribe and tax rates are expressed as fractions of sales, not profits. Since profits are much smaller than sales, the implied bribe and tax rates on profits are much higher than those on sales, so the estimated impact of a 1 percentage point increase in a tax on profits would be substantially smaller than what they estimate. Nevertheless, the large magnitudes they estimate remain something of a puzzle.

Another way that leakages could have efficiency consequences is through their impact on government provisions of goods and services. This could happen in two ways. First, if leakage simply increases the cost of government goods and services, this could have an effect similar to raising the price of these goods and services. The efficiency loss would arise if projects that would be cost effective at the true costs are no longer cost effective once the costs of corruption are included, and hence are not done. Second, leakages could create additional efficiency costs through distortions. Corrupt officials usually cannot steal cash directly, as that would be easily detected; instead, they need to go through a variety of more convoluted procedures to extract rents. These convoluted procedures themselves may induce inefficiencies, which could potentially be larger than the direct cost of corruption itself. We explore both of these issues in turn.

(i) *Price Effects*

One way corruption may matter is if theft of government resources increases the cost of government activity, so that otherwise worthwhile government projects—such as redistribution schemes or public works projects—become non-cost effective. Olken (2006) examines this possibility in the context of a large Indonesian anti-poverty program that distributed subsidized rice to poor households. As described above, by comparing survey data to administrative data, Olken estimates that at least 18 percent of the rice was lost from the program.

To examine the costs of corruption here, Olken performs a welfare calculation of the benefits of the program, both as it was implemented and using a counterfactual with the same targeting of beneficiaries but without corruption. The estimates imply that the welfare losses from this “missing rice” may have been large enough to offset the potential welfare gains from the redistributive intent of the program, so that the program without corruption might have been cost effective but, in the presence of corruption, it likely was not. In this particular case, the government implemented the program anyway, so in a sense the efficiency costs from lost redistribution were not realized. However, the estimates imply that in many similar cases, governments may not find redistribution programs worthwhile given the costs of corruption.

The Olken study demonstrates that a maximizing government might not have wanted to pursue the rice program given the costs of corruption in the program. The open question, however, is whether governments do indeed endogenously adjust their composition of expenditures in response to the higher prices imposed by corruption. To demonstrate these types of inefficiencies, it would be useful to show that if there was a shock to corruption that affected some sectors but not others, the government changed its composition of activities so that the quantity of activities where corruption declines increased. Note that this would be challenging to measure since we would be interested in measuring net quantities of corruption, not expenditures which include corruption—since if corruption on an activity declined, total expenditures on the activity (including corruption) could decline even though quantities could increase. We regard this question—of whether governments indeed optimize taking the price effects of corruption into account—as an important one for future research.

(ii) *Distortions*

Corruption may have another type of efficiency cost beyond the price effects discussed above. In particular, since corrupt officials need to hide their activity, they may introduce two types of distortions into the procurement of government activity. First, since corruption is secret, the government may not anticipate the amounts lost to corruption (in some ways, this is the countervailing force to the price effects discussed above). It may then effectively underfund some activities relative to its preferences, once the losses due to corruption are taken into account. Second, the need to keep corrupt activity secret could also introduce distortions, as procurement officials may substitute the types of goods that make hiding corruption easier. We discuss the evidence for both of these types of corruption in turn.

The first type of efficiency impact is the effective under-provision of government activities, since the government does not fully anticipate the impact of the losses due to corruption. As described above, Olken (2007) and Olken (2009) provide evidence for this type of efficiency loss in studies of perceptions vs. reality for rural roads in Indonesia. Since villagers are better able to detect corruption where prices are marked up (where there would only be a price effect), village officials instead hide their corruption by deflating quantities, i.e., they claim to procure enough rock, sand, and gravel to make a road that is 20cm thick but instead build a road that is only 15cm thick. Since the roads they build are thinner than official engineering guidelines, they will not last nearly as long, and will need to be replaced sooner. Although Olken was not able to directly detect this quicker rate of decay in the timeframe of his study, engineers estimate that the impact of the thinner-than-design roads on road lifespan is substantial enough to cause significant efficiency losses. Using a randomized design, Olken shows that reducing corruption in the program increases the thickness of the roads and brings them closer to engineering guidelines, reducing the efficiency losses.

Ferraz, Finan, and Moreira (2012) provide more direct evidence of the efficiency costs of corruption. They show that students in municipalities where corruption was detected in education have test scores that are 0.35 standard deviations lower than those without corruption, as well as higher dropout and failure rates. They demonstrate that higher corruption translates into lower quantities received: teachers in corrupt municipalities are 10.7 percentage points less likely to receive pedagogical training and less likely to have a computer or science lab. The study does not discuss the composition of school budgets, so it is hard to know if what the authors are picking up is price effects (there is less spending on schools because the government anticipates corruption) or distortions from corruption. One challenge in the study is that the level of corruption may be endogenous: while the authors control for other municipal characteristics, as well as corruption in other sectors and some indicators for school management practices, the level of corruption could be correlated with unobservable variables related to the quality of the school.

Examining leakages more broadly, Duflo, Hanna, and Ryan (2012,b) examine the impact of teachers' lack of attendance at schools. As described above, teacher attendance is a problem because it reduces a potentially key input of the learning process. In their study, their change incentives—coupled with monitoring—obtained a 21 percentage point lower rate of absenteeism, which led to a 0.17 standard deviation increase in student test scores. The results show that the lack of attendance was indeed translating into real education losses.

Another direct estimate of the efficiency costs due to distortion is the allocation of capital from state banks. Khwaja and Mian (2005) investigate whether politically connected firms, defined as those with a politician on their boards, obtain preferential lending. The results suggest politically connected firms

receive 45 percent larger loans from government banks in spite of having a 50 percent higher default rates on these loans. Privately owned banks, on the other hand, show no such political bias. According to estimates, and assuming the default rates are equivalent to transfers from taxpayers, the deadweight loss due to corrupt lending is between 0.15 percent and 0.30 percent of GDP. When the effect of inefficient investment of politically connected firms is considered, an additional 1.6 percent of GDP is estimated to be lost each year due to preferential lending.

Sukhtankar (2012) examines distortions in prices paid to farmers in sugar mills of Maharashtra, India during election years. He finds that in politically controlled mills, those that have a chairman who sought state or national office at least once between 1993 and 2005, farmers face lower prices by approximately Rs. 20 per ton during election years in comparison to prices paid to farmers in uncontrolled mills. In the year following the election, however, farmers in politically controlled mills receive higher prices than usual by about Rs. 80 per ton if their chairman wins the election. He finds that lower prices in election years reflect revenues which are embezzled for campaign expenses, but supporters are later rewarded. Sukhtankar's rough estimates indicate that a farmer can expect to receive about Rs. 27 per ton on their principal of Rs. 20 that is 'stolen' during an election year. As farmers are likely to be aware of this pattern of decrease and subsequent increase in sugar prices during and after an election, the author notes the siphoning off of revenues for campaigning is better thought of as rents paid to politicians rather than outright theft.

3.2.4. Impact on Correcting Externalities

A third way in which corruption may lead to inefficiency is if it lessens the government's ability to correct an externality. For example, if someone can bribe a police officer or judge instead of paying an official fine, the marginal cost of breaking the law is reduced from the official fine to the amount of the bribe. Even worse, if the police officer extracts the same bribe regardless of whether the person has broken the law, the marginal cost of breaking the law falls to zero and the law ceases to have a disincentive effect altogether.

Olken and Barron (2007) examined this possibility in their study of trucking in Aceh. Specifically, they examined what happens when these trucks stop at weigh stations. Overweight trucks are a classic example of an externality: the benefits to a trucker from loading on additional weight are concave, whereas the damage the truck does to the road rises to the 4th power with the truck's weight. For this reason, governments around the world weigh trucks and impose fines on trucks that are overweight.

Olken and Barron found that virtually all the trucks in our sample were substantially over the weight limits—and in fact, 42 percent of trucks were more than 50 percent over the legal weight limit. The data suggest that corruption at weigh stations is the likely culprit. Whereas according to the law all trucks more than 5 percent over the legal weight limit are supposed to be ticketed, immediately unload their excess cargo, and appear in court to face a fine, in fact virtually none of the trucks received an official ticket. Instead, virtually all paid a bribe. While more overweight trucks did pay higher bribes, this relationship was very flat, and even those trucks that were not overweight at all still had to pay a bribe. Corruption thus dramatically reduced the marginal cost of driving overweight, leading to more overweight trucks.

Bertrand et al. (2007) examined a similar question in the context of drivers' licenses in India. Specifically, they randomly allocated applicants for driving licenses into three groups. The first group received a bonus if they obtained a driver's license quickly, the second group received free driving lessons and the third group served as the comparison group. The findings confirm an efficiency loss from corruption: many people who were completely unable to drive were able to obtain licenses by paying a fee to an agent—and, in fact, the fee charged by the agent was unrelated to one's ability to drive. This efficiency loss effect was greater among the group that received the bonus for quickly obtaining a driving license since they faced a higher incentive to bypass the official procedures. On the other hand, those who were randomly allocated to the driver's license training class and who were better drivers were able to obtain their license with lower payments on average, mostly because they avoided using agents and instead used the official channel.

The Bertrand et al. (2007) and the Olken and Barron (2007) studies have very similar findings: in both cases, those who are doing the activity the government wishes to discourage (getting a license if you can't drive or having a truck that is overweight) do pay a higher cost than those who obey the laws. However, the marginal cost of breaking the law is much lower with corruption than it would be without corruption, so the net impact of corruption is to decrease the marginal cost of breaking the law and, thus, to decrease the effectiveness of the law.

Presenting a complimentary finding, Duflo et al. (2013) show that innovations that lead to decreased corruption can increase the effectiveness of laws and regulations. In their randomized field experiment of industrial plants in Gujarat, India, half of the plants continue to follow the status quo of hiring and paying their own third party auditors, and the other half are assigned to auditors with fixed pay from a central source whose work is subject to random verification checks by an independent technical agency. In comparison to the control group, they find that auditors in the treatment are 80 percent less likely to falsely report a plant's compliance with pollution regulations. Furthermore, knowledge of this increased accuracy in reporting actually compelled plants in the treatment group to pollute a significant .21 standard deviations less on average than those in the control group. Pollution reductions were entirely concentrated among the highest polluting plants, or those who face the largest risk of penalty.

These studies both raise an important question: given that corruption exists, how should the government structure the official laws so that the net of corruption marginal cost faced by citizens matches the government's true objective function? Put another way, if a sophisticated government knew there would be corruption at weigh stations, could they re-set the underlying fine schedule—which, after all, is the off-equilibrium punishment that determines the bribes that weigh station officials can extract—so that the net-of-corruption bribe schedule yields exactly the marginal cost curve they want? Studying this question—which would entail watching how corruption changes as the government changes the official cost schedule—is an important topic for future research in this area. It would also be informative to have evidence on this topic in sectors where the government is trying to correct externalities other than the transportation sector, such as criminal justice and environmental regulation.

3.2.5. Impact on Individuals

The final question is how corruption and leakages affect individuals. In part, the impact on individuals comes from all of the impacts described above: through the impact on firm behavior, through the efficacy of government service provision, and through the government's ability to correct externalities. But corruption and leakages can also affect individuals with the impact disproportionately distributed across society.

Hunt (2007) shows the negative distributional impact of corruption not by arguing that poor people expend a higher proportion of their income on bribes, but by stating that corruption can be an additional cost on the victims of misfortune—particularly crime victims. The study relies on an individual survey in Peru to show that misfortune increases victims' demand for public services, raising bribery indirectly. However, the study also shows that in many situations crime victims bribe more than other users who are not victims. Given the expense and disutility related to bribing, more corrupt governments would tend to aggravate the original problem faced by individuals in cases where they are vulnerable or more desperate for services.

In another investigation of the social impacts of corruption, Fisman and Wang (2013) compare worker fatality rates between politically connected and unconnected publically traded Chinese companies, using firm level data on worker fatalities from 2008-2011. They define a company to be politically connected if its CEO or another C-suite executive previously held a high level government position. They find that the rate of workplace deaths at connected firms is 5 times higher than at unconnected firms and that this relationship is best explained by firms using political connections to bypass safety regulations. Interestingly, however, when officials' promotion is contingent on meeting safety standards, then officials of connected firms are less prone to circumvent regulations, providing evidence that stricter penalties can lead to less abuse of power. Regardless, the study reveals that corruption can significantly impact individuals; in this case it leads to the disproportionate loss of life of workers whose firms ignore safety procedures due to their political connections.

3.2.6. Concluding Thoughts on Efficiency

The evidence in this section has demonstrated that the impacts of corruption and leakages on efficiency can be substantial. But it has also raised a number of important questions for new research.

One common theme that has emerged is that we know little about how governments respond endogenously to the presence of corruption. For example, if there are higher or lower rates of corruption in certain types of government spending, does it re-optimize spending as theory would predict, and does this re-optimization mitigate the efficiency costs of corruption? Or given that government rules to correct externalities are partially (but not completely) undone by corruption, does the government set official fines higher than they really want, knowing the official fines will not be implemented exactly? Understanding how the government responds to corruption is crucial to helping governments mitigate the efficiency losses from corruption in cases where it cannot be eliminated.

An issue on the flip side of this is the degree to which governments create regulations to maximize opportunities for corruption. An old hypothesis about corruption is that it actually increases efficiency, in

the sense that it allows citizens to cut through red tape and allows bureaucracies to be more responsive to economic needs. The Bertrand paper provided suggestive evidence of that view, in that people who had a high need to get a license quickly were able to do so by using an agent. But the red tape itself may be endogenously created as a way to maximize the corrupt rents captured by bureaucrats, as suggested by Banerjee (1997). Understanding whether the red tape itself is an endogenous response is an important, unanswered question, and another form of inefficiency that merits further study.

Similar issues apply to the costs of corruption for firms. While the Fisman and Svenson (2007) study suggested that bribes were more costly for firms than equivalent amounts of taxes, the tax rate could also be endogenous to the level of corruption. Gordon and Li (2009), for example, suggest that the tax code of developing countries is endogenously shaped by the presence of tax evasion, as governments reallocate tax systems towards those areas that are less prone to corruption. However, whether marginal tax rates on firms are higher or lower in corrupt countries and therefore whether the net distortions taxes imply for firms is higher or lower in corrupt countries, is an open question for future research.

Open Research Questions:

- *What are the macro impacts of corruption?*
- *What is the impact of corruption on a firm's performance, and why is the impact of corruption on firms higher than the impact of taxes?*
- *Does corruption cause governments to re-optimize expenditures? If so, do governments shift into lower corruption sectors (to avoid distortions) or higher corruption sectors (to maximize rents)?*
- *How do the bribes people have to pay to law enforcement depend on the official penalty functions? Given this knowledge, how can the government design the official law to best achieve its desired ends in the presence of corruption?*
- *Does the nature of corruption and its impact vary across different socioeconomic groups?*