Corruption in Developing Countries

Benjamin A. Olken[1] and Rohini Pande[2]

[1]Department of Economics, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142; email: bolken@mit.edu

[2]Harvard Kennedy School of Government, Harvard University, Cambridge, Massachusetts 02138; email: rohini_pande@harvard.edu

Keywords

Abstract

Recent years have seen a remarkable expansion in economists' ability to measure corruption. This in turn has led to a new generation of well-identified, microeconomic studies. We review the evidence on corruption in developing countries in light of these recent advances, focusing on three questions: how much corruption is there, what are the efficiency consequences of corruption, and what determines the level of corruption? We find robust evidence that corruption responds to standard economic incentive theory but also that the effects of anticorruption policies often attenuate as officials find alternate strategies to pursue rents.

1. INTRODUCTION

The past decade has seen a significant increase in the international policy community's interest in corruption. From 1998 to the present, 38 countries have ratified the OECD (Organization for Economic Co-operation and Development) Anti-Bribery Convention. At the end of 2005, the United Nations convention against corruption, the most comprehensive corruption convention to date, entered into force. In 2007, the World Bank launched its Strengthening World Bank Group Engagement on Governance and Anticorruption strategy. In recent years, the US Department of Justice and the Securities and Exchange Commission have dramatically increased their enforcements under the Foreign Corrupt Practices Act.[1] At the same time, several international aid agencies, including the Millennium Challenge Corporation, have made aid disbursements to low-income countries conditional on a country's corruption record.

These initiatives reflect a growing academic and policy consensus that corruption is often high in low-income countries and that it is costly. The growing policy activism that conditions international assistance on corruption outcomes in turn reflects a belief that given the right incentives, politicians, bureaucrats, and civil society in these countries can reduce corruption.

Evaluating these claims requires answers to three questions. The first, and arguably most basic, question that underlies policy design is, how prevalent is corruption? Second, what are the costs of corruption (i.e., is corruption actually harmful)? Finally, what factors influence the level of corrupt behavior? For example, is corrupt behavior responsive to economic incentives and market forces, and in what ways? Are there other effective approaches that can be brought to bear on corruption, such as technology, and how might corrupt officials react to such changes?

In this article we review the literature on these three issues and in each case describe both what we know and what we do not. We include a wide variety of types of corruption in our analysis but primarily focus on bribes to government officials and theft of government resources by public officials.[2]

In writing this review, several themes emerged. First, although there has been a revolution in the measurement of corruption over the past few years, estimated levels of corruption are remarkably heterogeneous, so there remains little consensus about its magnitude. Second, corrupt behavior has significant adverse consequences for efficiency and equity outcomes. Third, we find fairly robust evidence that corrupt behavior can be modeled in line with a few general economic principles: Corrupt officials respond to monitoring and punishments as one would expect from basic incentive theory, and standard market forces influence the level of bribes. That said, the ability of corrupt officials to substitute to alternate forms of corruption and to otherwise adapt to policy changes, either in the short run or the long run, suggests that applications of these principles can be tricky in practice.

In the end we were left with two very different senses. On the one hand, there has been a revolution in the measurement of corruption, and this has in turn led to a blossoming of

---

[1]In 2005 alone, the average number of new Department of Justice prosecutions exceeded fourfold the average for the prior five years.

[2]Although many laws (such as the US Foreign Corrupt Practices Act) also define cases of payment to any third party, in connection with or in furtherance of a contract, as corrupt practices, we exclude such cases from our analysis. Our definition of corruption also excludes shirking by government employees (here, employees steal time from the government, rather than money). A number of recent papers estimate the absenteeism of health workers and school teachers. We refer interested readers to Chaudhury et al. (2006) and Banerjee et al. (2012) for comprehensive reviews of that topic.

the academic literature on corruption. On the other hand, if we were asked by a politician seeking to make his country eligible for Millennium Challenge aid or the head of an anti-corruption agency what guidance the economic literature could give them about how to tackle the problem, we realized that, beyond a few core economic principles, we had more questions to pose than concrete answers.

Our review, especially the discussion of how to measure corruption, is related to recent survey articles, prominent among which are Zitzewitz (2012) and Banerjee et al. (2012). We complement these reviews by providing a summary of the different estimates of corruption magnitudes and identifying how anticorruption policies can be enriched by an understanding of the role of incentives and market forces in influencing corrupt behavior.

The remainder of this article proceeds as follows. Section 2 begins by reviewing the evidence to date on the magnitudes and efficiency costs of corruption. Section 3 begins by laying out a simple theoretical framework for thinking about the determinants of corruption and then examines evidence on them. Section 4 examines how corruption adapts to anticorruption policies. Section 5 concludes.

2. MAGNITUDES AND EFFICIENCY COSTS

Anecdotal and survey evidence suggests that corruption is rampant in the developing world and is more prevalent in developing countries than in rich ones (for a summary of the survey evidence, see Svensson 2005). Yet, as we show in Section 2.1, there are remarkably few reliable estimates of the actual magnitude of corruption, and those that exist reveal a high level of heterogeneity.

Just knowing the magnitude of corruption does not tell us how serious the problem is. After all, it is at least theoretically possible that corruption represents a transfer from one party (e.g., the government) to another party (e.g., bureaucrats), with little efficiency cost. In fact, if bureaucrats' official salaries were less than their market wage in expectation of the corrupt rents they would obtain—and there is evidence that this is indeed exactly what happens—there could be no net costs of corruption at all. In practice, however, the evidence we review in Section 2.2 suggests that the efficiency costs of corruption can be quite severe, as corruption may raise the marginal tax rate of firms, decrease business activity, raise the marginal costs of public funds, make certain government projects economically unviable, and undo the government's ability to correct externalities, leading to inefficient outcomes.

2.1. Estimating the Magnitude of Corruption

2.1.1. Perceptions. Until very recently, most estimates of corruption were based on surveys of perception. These perception surveys have the advantage of good coverage—it is much easier to ask someone's perceptions of corruption than to actually measure corruption directly. As such, they still form the basis of most cross-country corruption indices, such as Transparency International's Annual Corruption Perception Index and the World Bank's Control of Corruption Index.[3] Perception-based measures were also used in some

---

[3]The latter incorporates a variety of different aspects of corruption, ranging from the frequency with which firms make "additional payments to get things done," to the effects of corruption on the business environment, to measurements of "grand corruption" in the political arena.

of the first empirical work in economics on corruption, such as Mauro's (1995) cross-country study of the relationship between corruption and growth.

The challenge with perception-based measures is that they may not measure corruption accurately. To examine the reliability of villagers' perceptions of the corruption level in a local road building project, Olken (2009) obtains villager assessments of the likelihood of corruption in the project. At the same time, he develops a much more detailed measure of the amount of corruption that was actually present in the road project by comparing the amount the village government spent on the road to the amount independent engineers estimated the road would actually cost to build (for details, see Section 2.1.4). Although villagers' perceptions do reflect actual corruption in the road project, the correlation is quite weak: Increasing the actual missing expenditures in the road project by 10% increases the probability a villager reports any corruption in the road project by just 0.8%.

Moreover, villagers' perceptions appear to be biased in two ways. First, villagers are much better at detecting marked-up prices (i.e., overcharging for cement) than inflated quantities (i.e., billing for 1,000 m^3 of rocks but only delivering 800 m^3). To the extent that elected leaders care about villager perceptions, it is not surprising that most of the corruption occurs by inflating quantities. This may account for the relatively low correlation between perceptions and actual corruption, as people must make an inference about the aspects of corruption they cannot perceive—which end up being where the bulk of corruption is usually hidden. Second, Olken shows that individual characteristics, such as one's education, have much more power in predicting perceived corruption than actual corruption itself. If a perception survey has different compositions of respondents evaluating different projects (or countries), this could create systematic biases in the use of perception.

One response is to use expert surveys. Banerjee & Pande (2009) estimate political corruption among candidates for political office by surveying journalists who covered that election and politicians who stood for election in neighboring jurisdictions. They then correlate the reported outcomes (such as whether the candidate faced criminal charges) with actual data on the same and find a high correlation. The constraint on such surveys, however, remains researchers' ability to identify the correct expert pool, and of course in other settings it is possible that even experts' perceptions may be biased.

These types of biases could create problems in macro-level perception indices as well. For example, after the fall of Soeharto in 1998, many commentators perceived that corruption in Indonesia became worse. The commonly stated view was that many players at both the local and the national level started demanding bribes, and their failure to coordinate their bribe-taking behavior resulted in a higher total level of bribes. The worsening perceptions of corruption was captured by the Transparency International index—measured on a scale from zero (highly corrupt) to 10 (highly clean)—which fell from a value of 2.0 in 1998 to 1.7 in 1999, and stayed at the same level in 2000. This may well have been the case, but another explanation is that the fall of Soeharto's dictatorship resulted in a much freer press, which was newly able to report on allegations of corruption, which it did. It is therefore possible that perceptions of corruption rose even though actual corruption fell. For these types of reasons, economists have been moving to more direct measures of corruption whenever possible.

2.1.2. Survey estimates of bribes. Perhaps the most direct way of measuring bribery is through surveys of bribe payers. In most contexts, there is relatively little stigma associated with paying bribes, so in many cases bribery can be measured using surveys of firms or

households. One notable example is Svensson (2003), who surveys firms in Uganda and examines how much they paid in bribes. On average, firms in the survey report bribe payments of about USD 88 per worker, or approximately 8% of their total costs.

As this type of survey-based measure of bribes is the most easily replicable, it is one of the only areas where consistent measurement is now being carried out across countries and over time. One key data set is the International Crime Victims Survey from 49 countries, in which individuals are asked whether any government official in that country has asked them or expected them to pay a bribe for his services during the previous year. Using these data, Mocan (2008), for example, finds that the income and education of the individual have positive impacts on the likelihood of being asked for a bribe in developing, but not developed, countries. For firms, the World Bank Enterprise Surveys[4] have included comparable questions about firms' informal gifts or payments in obtaining water, electricity, telephone connection, operating and import licenses, and construction-related contracts; meeting with tax officials; securing government contracts; and more generally "getting things done" for many low- and middle-income economies. As this type of data becomes more available, we will be able to produce more reliable estimates of bribery over time and across countries.

2.1.3. Estimates from direct observation. The best way to measure corruption is often to observe it directly. Needless to say, this is difficult as officials rarely will let corrupt behavior be observed. Nevertheless, there are several notable examples of the direct observation of corrupt activity. One such example is the case of Montesinos in Peru, documented by McMillan & Zoido (2004). Montesinos, who was secret-police chief under President Alberto Fujimori in Peru, bribed judges, politicians, and the news media to support the Fujimori regime. Remarkably, he kept detailed records, with signed contracts from those he bribed and videotapes of them accepting the bribes, and these became public after the fall of the Fujimori regime. McMillan & Zoido use them to estimate the cost of bribing various types of government officials. On average, politicians received bribes ranging from USD 3,000 to 50,000 per month, depending on whether the politician was in the opposition party (higher) or Fujimori's party (lower), with judges receiving bribes of the same order of magnitude. The bribes to control the media were orders of magnitude larger—as much as USD 1.5 million per month for one television station's support.

Olken & Barron (2009) provide direct data on actual bribes in a more prosaic setting: the bribes truck drivers pay to police on their routes to and from the Indonesian province of Aceh. Over a nine-month period, enumerators accompanied truck drivers on their regular routes, dressed as truck drivers' assistants, and simply noted the amounts that truck drivers paid each time they were stopped at a police checkpoint or weigh station. On over 300 trips, they observed more than 6,000 illegal payments. Usually each payment was small, averaging USD 0.50 to 1, sometimes in cash and sometimes in kind (such as a pack or two of cigarettes). In total, the illegal payments represented 13% of the marginal cost of the trip. By comparison, the salary of the truck driver was only 10% of the marginal cost of the trip.[5]

---

[4]The exact details on the number of countries and years for each type of survey are available at **https://www. enterprisesurveys.org**.

[5]The authors also compared directly observed bribes to reported bribes from a survey of comparable trips and found that reported bribes were about double the amount of actually observed bribes. One potential explanation is that drivers have an incentive to overreport bribes in general because they are reimbursed by trucking firms on the basis of the average amount of bribes they need to pay.

Sequeira & Djankov (2010) use a similar methodology in Mozambique and South Africa, shadowing clearing agents who process customs for cargo as it passes through the ports. Specifically, they estimate the economic costs and distortions associated with 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. They find that, on average, bribes represent 14% of the shipping costs for a standard container passing through the port of Maputo, Mozambique, and 4% of the shipping costs for a standard container passing through Durban, South Africa.

2.1.4. Graft estimation by subtraction. The most common method for estimating graft (i.e., the theft of government funds) is what we term estimation by subtraction. In this method, one obtains two measures of the same quantity: one measure before corruption takes place and one measure after corruption takes place. The estimate of corruption is the difference between the two measures.

One of the first estimates using this technique is the pioneering study by Reinikka & Svensson (2004). Using what they term a public expenditure tracking survey (PETS), they compare the amount of a special education block grant sent down from the central government in Uganda with the amount of the block grant received by schools. They estimate a leakage rate of 87%. Once the results were publicized, a subsequent study found that the leakage rate fell to less than 20%. An important question in such an approach is the quality of record keeping: If schools have poor records, some of the money might not show up on the books even though it may have been received. Studying the importance of record-keeping quality in PETS is an important issue for the replicability of this technique. Subsequent to this work, similar PETS studies have been carried out, largely by the World Bank, in a variety of contexts (for a brief review, see Olken & Pande 2011).

Using a similar approach, Fisman & Wei (2004) measure tax evasion by comparing Hong Kong's reported exports and China's reported imports of the same products. They differentiate three different aspects of tax evasion: underreporting of unit value, underreporting of taxable quantities, and mislabeling of higher-taxed products as lower-taxed products. These calculations are then used to estimate the effect of tax rates in tax evasion. They found that higher-taxed products were associated with a 40% higher median evasion rate.

Olken (2007) implements a related exercise in the case of rural road projects. He compares the official amount spent on the road to an independent engineering estimate of what the road actually cost to build, in which engineers dug core samples of the roads to estimate materials quantities, did price surveys to estimate local prices, and interviewed villagers to estimate actual wages paid. Importantly, because some amount of materials naturally disappears during construction, Olken built several small test roads where he knew there was no corruption so that he could calibrate the metric so it would show zero corruption when, in fact, corruption was zero. Olken estimated that "missing expenditures"—the difference between what the village claimed the road cost and what the engineers estimated it actually cost—averaged approximately 24% of the total cost of the road.

An alternative approach is to compare administrative data to a generally administered household survey. Olken (2006) uses this approach to estimate theft of rice from a program that distributed subsidized rice in Indonesia. He estimates that, on average, at least 18% of the rice cannot be accounted for, with greater amounts in ethnically heterogeneous and sparsely populated areas. In a similar vein, Niehaus & Sukhtankar (2010) compare administrative and survey data to measure corruption as the gap between official and actual

quantities, including over-reporting of days and underpayment of wages in the Indian National Rural Employment Guarantee Act.

When examining corruption through price manipulations, one can compare an official price to the market price and use the difference as a measure of price manipulation. Hsieh & Moretti (2006) do this for a famous case: corruption under the Iraqi Oil-for-Food program administered by the United Nations. Specifically, they compare the price received by Iraq for its oil to the going price for comparable oil on the world spot market and use a model of the market for oil trading to infer what share of that underpricing was likely received by Saddam Hussein's regime. Although the total amount of corruption they estimate is enormous—approximately USD 1.3 billion—it amounts to only approximately 2% of the total volume of oil sold. Of course, not all price markups are corruption—they could simply reflect incompetence or a lack of incentives in obtaining good prices for the government (see, for instance, Bandiera et al. 2009).

2.1.5. Estimates from market inference. In some cases, one can use the theory of market equilibrium, combined with data on market activity, to estimate the amount of corruption. In a pioneering study, Fisman (2001) applies this approach to estimate the value of political connections to Indonesian President Soeharto. Specifically, he obtains an estimate from a Jakarta consulting firm of how much each publicly traded firm was "connected" to Soeharto, on a scale of zero to four. He then estimates how much each firm's price moved when Soeharto fell ill to estimate the stock market assessment of the value of those political connections. If the efficient-markets hypothesis holds, then the change in stock market value surrounding these events captures the value of the political connection to the firm. Because investment bankers in Jakarta estimated that the total market would fall by 20% if Soeharto died, he can calibrate these estimates to estimate the total "value" of the connections to Soeharto. On net, for the most connected firms, he estimates that approximately 23% of their value resulted from Soeharto's connections.

The Fisman market approach is replicable in any case in which one has data on firms' connections to prominent politicians and when the politician experiences health shocks. For example, Fisman et al. (2006) replicate the same approach for the United States, looking at the value of connections to former US Vice President Dick Cheney, using shocks while he was a candidate and while he was in office. In a marked contrast with the Soeharto paper, they find that Cheney's heart attacks had zero effect on the value of Cheney-connected stocks.

Faccio (2006) pursues a similar approach using a large sample of countries—she examines political connections to 20,202 publicly traded firms in 47 countries. For each of these firms, she defines the firm as having a political connection if a board member or large shareholder is a politician (e.g., member of parliament or minister). She focuses on corporations in which a previous board member or large shareholder becomes a politician. She finds that, on average, having a member of the board or large shareholder become a politician is associated with a 2.29% increase in the company's share value. Echoing the contrast between Soeharto in Indonesia and Cheney in the United States, when she splits the sample into countries with below- and above-average corruption levels (as measured by the World Bank perceptions index), she finds that the impact comes entirely from high-corruption countries: In above-median-corruption countries, having a board member or large shareholder become a politician increases stock market value by 4.32%, but in below-median-corruption countries, having a board member or large shareholder become a politician has no impact on stock value.

Another approach to measuring corruption uses equilibrium conditions in the labor market. Specifically, one can use the fact that people in the public sector must, on the margin, be indifferent between their public sector job and alternative jobs in the private sector. If, controlling for their job opportunities, pay is lower in the public sector, the result could simply reflect a compensating wage differential. But if pay in the public sector is lower but consumption levels are the same, one could infer that the difference between pay and consumption in the public sector relative to the same difference in the private sector tells us something about how much those in the public sector are likely receiving in the form of bribes. Gorodnichenko & Peter (2007) perform this exercise using a household survey in Ukraine. They find that, controlling for education, hours of work, job security, fringe benefits, job satisfaction, and secondary employment, public sector workers received 24%–32% less income than their private sector counterparts. Crucially, however, they have the same level of consumption and assets, suggesting that a large part of the gap must be made up in bribes. Aggregating across the economy, they estimate that the total amount of the gap (and hence bribery) is between USD 460 million and 580 million, or approximately 1% of GDP.

2.1.6. Other approaches. Although above we discuss the main approaches used in the literature, this is not an exhaustive list. For instance, Ferraz & Finan (2008, 2011) and Brollo et al. (2010) use official audits of municipal governments in Brazil to identify instances of corruption. These audits were directly summarized and made available to the media. The summary provided a short description of each irregularity in the municipality. The challenge with audit data is that they represent a combination of both actual corruption and the inability to hide it from auditors, so these data need to be used with care.

Asking whether, conditional on observables (which measure eligibility for public programs), public officials are more likely to benefit from publicly provided private transfers gives another measure of corruption. Besley et al. (2012) show that controlling for asset-based eligibility, holding political office increases the likelihood that a villager in India has a Below Poverty Line card by 10%. A similar approach (and findings) is reported by Olken (2007) and Niehaus et al. (2012).

2.1.7. So how much corruption is there, really? **Table 1** presents the magnitude of corruption estimated from all the studies reviewed above, separated into estimates of graft (theft of government funds) and estimates of bribes.[6] The table shows the dramatic range. It also shows that, although a number of credible estimates have emerged, in some sense there are relatively few hard data when compared with other development indicators.

The magnitudes of corruption raise several important questions. First, a striking correlation that comes up in a variety of data sets—from the perception indices, to Faccio's (2006) and Fisman's (2001) studies of the value of political connections, to Sequeira & Djankov's (2010) comparison between ports in South Africa and Mozambique—is the strong negative relationship between income and corruption: As best we can measure it, richer countries appear less corrupt. The causality potentially runs in both directions. It is easy to see how low corruption could cause countries to become rich if corruption hinders

---

[6]We include estimates of the value of political connections in the graft category, under the idea that the value of those connections comes from the firm's ability to appropriate rents from the government due their connections, although one could easily categorize them separately instead.

Table 1 Magnitudes of corruption

Paper	Country	Context	Strategy for assessing corruption	Corruption estimate	Corruption estimate (%)
Estimates of bribery					
Svensson (2003)	Uganda	Bribes firms paid	Survey evidence	Firms paid bribes of USD 88 per worker	8% of costs
Olken & Barron (2009)	Indonesia	Bribes truck drivers paid to police on their routes	Direct observation; enumerators accompanied truck drivers on their regular routes, dressed as truck drivers' assistants, and observed illegal payments	Truck drivers paid bribes averaging USD 0.50 to 1 per payment	13% of cost of a trip
McMillan & Zoido (2004)	Peru	Bribes the secret police paid to judges, politicians, and media to support the Fujimori regime	Direct observation; after fall of Fujimori regime, videotapes and bribe receipts became public	Politicians received bribes of USD 3,000–50,000 per month; media received bribes as much as USD 1.5 million per month for one television station's support	N/A
Sequeira & Djankov (2010)	South Africa	Bribes paid to port and border post officials	Direct observation; enumerators shadowed clearing agents in ports to collect information on bribe payments	Bribes amounted to 14% and 4%, respectively, of the total shipping costs for container passing through Mozambique and South Africa	14% of shipping costs (Moz.)
	Mozambique				4% of shipping costs (S. Africa)
Estimates of graft					
Reinikka & Svensson (2004)	Uganda	Graft in public spending of educational funds intended to cover schools' nonwage payments	Estimate by subtraction; PETS compared the amount of grant sent down from the central government to the amount received by schools	Schools received on average only 13% of the grants	87% of funds

(*Continued*)

Table 1   (*Continued*)

Paper	Country	Context	Strategy for assessing corruption	Corruption estimate	Corruption estimate (%)
Estimates of graft					
Olken (2007)	Indonesia	Graft in the building of rural roads funded through a national government program	Estimate by subtraction; the official amount spent on the road was compared to an independent engineering estimate of what the road actually cost to build	"Missing expenditures" (the difference between what the village claimed the road cost and what the engineers estimated it actually cost) averaged approximately 24% of the total cost of the road	24% of cost of the road
Olken (2006)	Indonesia	Theft of rice from a program that distributed subsidized rice	Estimate by subtraction; administrative data were compared to a generally administered household survey	At least 18% of the program's rice disappeared before reaching households	18% of program expenditures
Hsieh & Moretti (2006)	Iraq	Bribes from the underpricing of oil in Iraq's Oil-for-Food program	Estimates by subtraction; the selling price of Iraqi oil was compared to the Oil-for-Food program and the author's estimates of the "market" price of Iraqi oil	Iraq collected USD 1.3 billion in bribes from underpricing oil, or 2% of oil revenues	2% of oil revenues
Khwaja & Mian (2005)	Pakistan	Politically connected loans	Estimate from market inference; the additional nonperformance rate from politically connected loans was compared to that from nonpolitically connected loans	Politically connected firms received 45% larger loans from government banks despite having 50% higher default rates	0.3%–1.9% of GDP

(*Continued*)

Table 1 (*Continued*)

Paper	Country	Context	Strategy for assessing corruption	Corruption estimate	Corruption estimate (%)
Estimates of graft					
Niehaus & Sukhtankar (2010)	India	Wages on the National Rural Employee Guarantee Scheme	Estimate by subtraction; the officially paid wages were compared to wages as reported by a survey	INR 236 were stolen per actual day paid, where an actual day paid is about INR 60	79% of labor expenditures
Fisman (2001)	Indonesia	Value of political connections to President Soeharto for Indonesian public firms	Estimate from market inference; firm stock price movement was examined when Soeharto fell ill, given the strength of its political connections to Soeharto	For the value of the most connected firms, 23% was due to political connections	23% of firm value
Fisman et al. (2006)	United States	Value of personal ties to Vice President Cheney for US public firms	Estimate from market inference; firms' stock price movement was analyzed in response to the shock to Cheney's health, given the strength of connection to Cheney	In all events studied, there is zero effect on the stock prices of connected firms	0% of firm value
Faccio (2006)	Cross-country	Value of political connections for firms across sample of 47 countries	Estimate from market inference; stock price movements of firms were analyzed around the time of announcements that officers or large shareholders were entering politics or that politicians were joining their boards	There is a 2.29% increase in company value when a businessman enters politics and a 4.32% increase in stock market value when a board member or large shareholder becomes a politician in countries where corruption is above the median	2.3%–4.3% of company value

Table 1 (*Continued*)

Paper	Country	Context	Strategy for assessing corruption	Corruption estimate	Corruption estimate (%)
Estimates of graft					
Gorodnichenko & Peter (2007)	Ukraine	Bribes received by public sector employees	Estimate from market inference; residual wage differentials between the public and private sectors were examined (consumption levels are the same in the two groups, and labor market equilibrium implies that employees are indifferent between working in the public and private sectors)	The aggregate amount of bribery was estimated to be between USD 460 million and 580 million, or approximately 1% of the GDP of Ukraine	1% of GDP
Ferraz & Finan (2010)	Brazil	Corruption in municipal government	Audit reports from central government auditors	Audits find an average of BRL 327,000 diverted resources per violation, or 8% of total amount audited	8% of total amount audited
Besley et al. (2012)	India	Beneficiary selection by village council	Conditional on eligibility, does political office predict beneficiary status	Chief village councilor is 10% more likely to be a beneficiary	2% of beneficiaries selected

economic activity (Mauro 1995). However, the relationship in the other direction—that richer countries become less corrupt—is less obvious. On the one hand, certain types of income shocks, such as natural resource shocks, may lead to there being more rents to be expropriated and more corruption. For example, Caselli & Michaels (2009) present the case of oil revenues distributed to municipalities in Brazil, as a result of the large increase in Brazil's offshore oil production, and argue that this led to an increase in corruption. There is, however, some evidence that these rents dissipate in the medium run, possibly because voters become more aware of the total resources (Monteiro & Ferraz 2010). On the other hand, more complex business relationships may lead to demand for better government, and higher incomes may mean that countries have more resources to invest in cleaning up corruption (Treisman 2000).

Second, even among countries at similar income levels, and even within countries, there is marked heterogeneity in corruption levels, as shown in **Table 1**. Once one starts

examining why corruption emerges, it becomes clear that there is no reason to expect magnitudes of corruption to be similar across settings.

Third, virtually all of these "hard" estimates of corruption may suffer from selection bias in both directions. To the extent that measures of corruption depend on voluntary disclosure (such as surveys of bribery or disclosing links to politicians sitting on corporate boards), corruption may be understated, as places where corruption is most severe might be less likely to disclose it. To the extent that researchers purposefully choose cases to study, corruption may be overstated, as researchers may hone in on situations where they expect to find corruption. Developing careful, rigorous metrics of corruption that are not subject to these types of selection bias is an important area for future research.

2.2. Does Corruption Matter?

Although the previous section shows that corruption is substantial in magnitude—whether in the form of bribes given to civil servants or graft from public expenditures—this does not necessarily answer the question of whether corruption actually has a negative impact on economic activity.

For example, Gorodnichenko & Peter (2007) show that, on average, public employees in Ukraine have the same consumption levels as their private sector counterparts, even though their salaries are 24%–32% lower. Corruption 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. In this case 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 eliminated. More generally, corruption could either have efficiency costs or lead to efficiency gains.

This section lays out the evidence thus far on the ways in which corruption may have aggregate efficiency costs: 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. The endogenous nature of corruption makes finding credible instruments for corruption at the macro level difficult. We therefore restrict attention to micro evidence.

2.2.1. Impact on firms. To estimate the efficiency cost of corruption on firm behavior, ideally one must know several things. First, how does corruption change 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), the marginal bribe rate should be below the official marginal tax rate, so corruption reduces effective tax rates. Conversely, if bribes are charged for other types of government activities, this could increase the effective marginal tax rate faced by firms. Second, conditional on knowing the effective marginal tax rate after corruption, for a given effective marginal tax rate, are taxes affected by corruption more distortionary than de jure taxes?

Svensson's (2003) study of bribe-paying behavior in Uganda provides some clues that although 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% 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% 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% on profits and 0.4% on capital, one might expect a modest impact of corruption on firm activity. By way of comparison, the marginal tax rate on corporate profits for large corporations in the United States is 35%.

Svensson's study establishes effective corruption tax rates but does not tell us the impact of corruption on firms. There may be other ways in which corruption affects 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) makes this argument looking at foreign direct investment and measuring uncertainty through perceptions-based metrics. More recently, Malesky & Samphantharak (2008) use survey data to show that changes in governors in Cambodia are associated with increases in uncertainty about corruption but with reductions in actual corruption levels and decreased firm-level investment.

In Section 2.1.3 we discuss Sequeira & Djankov (2010), who examine a different type of distortion: changes in the firm's production choices designed to avoid corruption. Their estimates suggest that approximately 46% of South African firms located in regions where overland costs to the port of Maputo are 57% 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 a change in the effective marginal tax rate) and indirect effects (through uncertainty or other channels), it is necessary to directly examine the net impact of corruption on firm decisions. Using the same data set as Svensson (2003), Fisman & Svensson (2007) calculate bribes and tax payments in Uganda as a function of total firm 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. A 1-percentage-point increase in bribes reduces annual firm growth by 3 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.[7]

2.2.2. Impact on government provision of goods and services. Corruption can have efficiency consequences through impacts on government provisions of goods and services. First, if it 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, corruption 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

---

[7]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. Because 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.

induce inefficiencies, which could potentially be larger than the direct cost of corruption itself. We explore both these issues in turn.

Price effects.   One way corruption may matter is if the 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 antipoverty program that distributed subsidized rice to poor households. As described above, by comparing survey data to administrative data, Olken estimates that at least 18% of the rice was lost from the program. He also 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. An open question, however, is whether governments endogenously adjust their composition of expenditures in response to the higher prices imposed by corruption. We regard this question—of whether governments indeed optimize taking the price effects of corruption into account—as important for future research.

Distortions.   Because corrupt officials need to hide their activity, they may introduce two types of distortions into the procurement of government activity. First, because 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 these types of corruption in turn.

The first type of efficiency impact is the effective underprovision of government activities, as the government does not fully anticipate the impact of the losses due to corruption. As described above, Olken (2007, 2009) provides evidence for this type of efficiency loss in studies of perceptions versus reality for rural roads in Indonesia. Because villagers are better able to detect corruption when prices are marked up (for which 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 20 cm thick but instead build a road that is only 15 cm thick. As 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 time frame of his study, engineers estimate that the impact of the thinner-than-designed roads on road life span is substantial enough to cause significant efficiency losses.

Ferraz et al. (2010) provide direct evidence of the efficiency costs. In Brazilian municipalities where corruption was detected in education, they show that students have test scores that are 0.35 standard deviations lower than those municipalities without corruption, as well as higher dropout and failure rates. Moreover, higher corruption translates into lower quantities received: Teachers in corrupt municipalities are 10.7 percentage points less likely

to receive pedagogical training and are 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 are 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: Although 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.

Another direct estimate of the efficiency costs due to distortion is the allocation of capital from state banks. Khwaja & Mian (2005) show that politically connected firms, defined as those with a politician on their boards, receive 45% larger loans from government banks despite having a 50% higher default rate on these loans. Privately owned banks, conversely, 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% and 0.30% of GDP. When the effect of the inefficient investment of politically connected firms is considered, an additional 1.6% of GDP is estimated to be lost each year due to preferential lending.

2.2.3. 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 & Barron (2009) examine this possibility in the context of trucks stopping at weigh stations in Aceh, Indonesia. 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 fourth power with the truck's weight. They found that almost all trucks were substantially over the weight limits—and 42% of trucks were more than 50% over the legal weight limit. The data suggest that corruption at weight stations is the likely culprit. Whereas all trucks more than 5% over the legal weight limit are supposed to be ticketed, immediately unload their excess cargo, and appear in court to face a fine according to the law, in fact virtually none received an official ticket. Instead, almost all paid a bribe. Although more overweight trucks did pay higher bribes, this relationship was very flat, and even those trucks that were not overweight had to pay a bribe. Corruption thus dramatically reduced the marginal cost of driving overweight, leading to more overweight trucks.

Bertrand et al. (2007) examine a similar question in the context of driver licenses in India. They randomly allocated applicants for driver 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: 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 because they faced a higher incentive to bypass the official procedures. Conversely, 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.

Bertrand et al.'s (2007) and Olken & Barron's (2009) studies have similar findings: In both cases, those who are doing the activity the government wishes to discourage (getting a license if they cannot drive or having a truck that is overweight) do pay a higher cost that 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.

These studies raise an important question: Given that corruption exists, how should the government structure the official laws so that the net of the corruption marginal cost faced by citizens matches the government's true objective function? Understanding how corruption maps de jure marginal costs imposed by laws into de facto marginal bribe payments (and hence the de facto marginal costs faced by individuals) is an important next step in thinking about how to more effectively write laws in the presence of corruption.

2.2.4. Impact on individuals. A final question is how corruption affects individuals directly. 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 for 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.

2.2.5. Some concluding thoughts on efficiency. 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 the government reoptimize spending as theory would predict, and does this reoptimization 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 it really wants, knowing the official fines will not be implemented exactly?

An issue on the flip side of this is the degree to which governments create regulations to maximize opportunities for corruption. A classic example is red tape. We refer the interested reader to Banerjee et al. (2012), who develop a theoretical framework for understanding how red tape itself may be endogenously created as a way to maximize the corrupt rents captured by bureaucrats, as suggested by Banerjee (1997), and review the existing empirical literature. It is clear that understanding whether the red tape itself is an endogenous response is another form of inefficiency that merits further study.

Similar issues apply to the costs of corruption for firms. Although Fisman & Svensson (2007) suggest that bribes are more costly for firms than the equivalent amounts of taxes, the tax rate could also be endogenous to the level of corruption. Gordon & 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 toward 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 that taxes imply for firms are higher or lower in corrupt countries, is an open question for future research.

3. WHAT DETERMINES CORRUPTION?

This section examines what we know about why corruption exists and, relatedly, what can be done with it. To organize ideas, we provide a simple framework that models the perspective of an individual bureaucrat, following the ideas of Becker & Stigler (1974). This framework treats the gains from corruption (the bribe) as fixed and asks when honesty will be preferable to dishonesty. We then examine what happens when the optimal bribe is determined by the bureaucrat taking into account market forces, following the ideas of Shleifer & Vishny (1993). The subsequent sections discuss the empirical evidence along the dimensions suggested by the simple theoretical framework.

3.1. The Incentives Bureaucrats Face

Suppose that the bureaucrat receives a wage w from the government and, if fired, can receive an outside option v. The bureaucrat can decide to be corrupt or honest. If corrupt, he is detected with probability p, is fired, and receives outside option v. If he is undetected, he receives his wage w plus the bribe b, less a dishonesty cost d. In equilibrium, he will be corrupt if and only if $w - v < \frac{1-p}{p}(b-d)$.

   This framework suggests several avenues for reducing corruption. One could increase the returns to staying on the job (w), or, equivalently in this context, one could decrease the outside option (v) by increasing punishments. One could also increase the probability of detection (p).

   One implication is that if there is heterogeneity in d among potential bureaucrats, there can be selection in which those who are most likely to be corrupt (those who have the lowest dishonesty costs d) will self-select to be more likely to become bureaucrats. Suppose that d in the population is distributed uniformly from zero to $\bar{d}$. If $w - v > \frac{1-p}{p}(b)$, then nobody will be corrupt, regardless of their level d, and there is no reason that the distribution of d among bureaucrats will be different than the distribution of d in the population.

   If, however, the above inequality does not hold, then people with low d will have a higher utility from becoming bureaucrats than those with high d, as they will be relatively more efficient at corruption. Therefore, depending on how the government allocates jobs, we might expect to have more low-d people among bureaucrats than in the population. This implies that corruption may be harder to combat because a corrupt system may attract bureaucrats who are more prone to corruption. It also implies that the effect of a given anticorruption policy [i.e., a vector (w,p)] will depend on past levels of anticorruption policies, as those past policies will influence the selection of bureaucrats.

   The simple framework thus far has treated the amount of the bribe, b, as exogenous. In practice, however, the bribe may be set by the bureaucrat to maximize his profits. Specifically, conditional on deciding to be corrupt, the bureaucrat will set his bribes to maximize his profits, which are the number of bribes he receives multiplied by the price (i.e., $q_i b_i$).[8] The key insight of Shleifer & Vishny (1993) is that the optimal solution depends on what other bureaucrats are doing and how they set prices. If a person needs permits from two

---

[8]To simplify the analysis, we assume that conditional on being caught, the probability of being caught does not depend on the amount of the bribes or the quantity of bribes, although one could easily generalize the model to include these effects.

different bureaucrats to complete a transaction, and both set prices independently, then each bureaucrat solves $\max_{b_i} Q(b_i + b_{-i})b_i$, taking the other bureaucrat's bribe b_{-i} as given. In such a case, the total amount of the bribes $(b_1 + b_2)$ will be higher than if there had only been a single bureaucrat, and the total quantity will be lower. Conversely, if a consumer needs a single permit that can be obtained from either bureaucrat, they will compete against each other and reduce the bribe beyond what a single, monopolistic bureaucrat would charge. The key insight is that the bribes themselves may be a function of the structure of the bureaucracy and that changing the nature of the organization may have important implications for the level of corruption.

This framework, although stylized, highlights the important role that both the incentive structure faced by individual bureaucrats (be it compensation, monitoring, selection, or other incentives) and the bureaucratic organization may play in influencing the amount of corrupt behavior. This section discusses the evidence to date on each of these factors in turn.

3.1.1. The bureaucrat's decision problem: the role of compensation, selection, and other incentives. We examine several aspects of the bureaucrat's decision problem: the role of compensation, monitoring and punishment, selection, and incentives.

Compensation. Despite the attention often given to civil service wages, there is relatively little evidence on their impact. Several cross-country studies find that higher public wages are associated with lower corruption, although these studies are essentially cross-sectional in nature. For instance, in a cross section of 31 low-income countries, Van Rijckeghem & Weder (2001) find that a doubling of government relative to manufacturing wages is associated with only a 0.5-point reduction in the International Country Risk Guide's corruption index measured on a scale from zero to six. Meanwhile, Rauch & Evans (2000) find that the level of bureaucratic wages is significant in explaining only one of the five measures of bureaucratic performance, namely that a 1-standard-deviation increase in salary is associated with an improvement of 0.5 standard deviations in the bureaucratic delay index measured on a scale from one to four.

With regard to more micro evidence, Di Tella & Schargrodsky (2004) test the efficiency wage idea by looking at a corruption crackdown in hospitals' procurement departments in Buenos Aires. They examine the impact of increasing the probability of detection and examine heterogeneous impacts on the prices paid for basic inputs based on the level of wages. Prices paid by hospitals for basic, homogeneous inputs decrease by 15% during the first nine months of the crackdown, and the following period prices increase, but remain 10% lower than those prevailing before the crackdown. During the first phase of the crackdown, when audit intensity can be expected to be maximal, higher wages have no effect on inducing lower input prices. Meanwhile, higher wages do have a negative effect in the last phase of the crackdown, when audit intensity can be expected to take intermediate values—the wage elasticity of input prices exceeds 0.2.

Niehaus & Sukhtankar (2010) examine the idea that the rents from keeping one's job can deter corruption today to preserve tomorrow's opportunities. The rents they examine come from corruption, not wages. They can identify the effect of future rents on the level of corruption today because the program features two types of projects; some projects pay fixed daily wages, whereas others pay piece rates. They examine how corruption in the two types of projects varies with anticipated rent-extraction opportunities using an exogenous increase in the wage rate for daily wage projects. Their results show an 80% reduction in

the daily theft on piece-rate projects in the period following the wage increase. Hence, when the opportunities for theft from daily wage projects increase, theft on piece-rate projects goes down. In addition, they find reduced over-reporting of days worked on daily wage projects in areas where the proportion of future daily wage projects is higher.

Monitoring and punishments. One would expect from the above framework that increased monitoring would reduce corruption. In practice, however, the very individuals tasked with monitoring and enforcing punishments may themselves be corruptible, so increasing monitoring may simply increase transfers from low-level officials to auditors. Moreover, just because people are audited does not necessarily mean that auditors will find enough evidence to actually impose a punishment, even if corruption were taking place. Understanding the degree to which additional monitoring can reduce corrupt behavior is thus an important area for empirical research.

In his study of roads in Indonesia, Olken (2007) examines this question by conducting a randomized experiment on auditing. Before villages began building road projects, some villages were randomly selected for a high-audit-intensity group, where they faced an audit by the government agency with 100% probability, as opposed to a 4% probability in the control group. Olken finds substantial effects of the government audits, reducing corruption by 8 percentage points or approximately 30% from the baseline level. Interestingly, the audits revealed substitution among alternative forms of corruption: Although audits reduced missing expenditures, they increased nepotism (i.e., the hiring of family members of the project leader or village officials to work on the project). One reason that the audits did not reduce corruption to zero was that, even though audits found problems in 90% of the villages they audited, the findings were typically administrative failures, such as improper receipts or a failure to receive the required number of competitive bids, rather than the direct evidence of corruption that would be needed for a criminal prosecution. In other words, just because the probability of an audit was 100% does not imply that the probability of punishment, conditional on the presence of corruption, was 100%. Nevertheless, on balance, the results demonstrate that the traditional economic approach to fighting crime—increasing the expected cost of crime by increasing the probability of being caught—can play an important role in reducing corruption, even in a highly corrupt environment where those doing the monitoring are themselves potentially corruptible.

Another approach to providing monitoring—that does not involve a central auditor—is grassroots monitoring, in which regular citizens are empowered to monitor their officials to prevent corruption. Olken's study also examines this by randomly allocating villages to receive more intensive community monitoring. This was done through two interventions, with different purposes. The first intervention involved inviting hundreds of villagers to attend local accountability meetings, to reduce elite control over which community members were involved in the monitoring. The second intervention involved distributing hundreds of anonymous comment forms throughout the village, to allow community members to voice concerns or complaints without fear of retaliation.

The invitations intervention reduced theft of materials, but only for theft of wages (i.e., convincing villagers to work for free but billing the project for their work). One reason may be that if theft of wages was detected, the benefits would go to the small number of people who worked on the project and should have been paid; they therefore have a strong personal incentive to prevent this type of corruption. By contrast, the benefits from detecting theft of materials would accrue to the village as a whole in the

form of a better road, so the free-rider problem may be more severe. With regard to the anonymous comment forms, they were successful only when they were distributed via school children, not via the neighborhood government, as the neighborhood leaders channeled the forms toward preferred people who were more likely to support the elite in the project. One important thing to take away from is that for community participation to work, it is important to get the details right in terms of protecting people from retaliation, limiting the free-rider problem, and preventing elite capture.

Increasing community participation can influence governance through multiple channels. The first, which is emphasized by Olken, is improvements in monitoring. A second possible channel is improved information—leaders may learn more about villagers' preferences, and villagers in turn may learn more about outcomes. The second channel may be particularly important when the outcomes being measured relate to service delivery, not corruption.[9]

Ferraz & Finan (2008) examine the role of electoral sanctions. Randomized timing of municipality audits allows them to examine the impact of audit timing on the probability that the mayor is re-elected. Conditioning on the number of corruption violations found by the auditors, those audited before the election were less likely to be re-elected than those who were audited after the election. The finding suggests an important complementarity between audits—which provide information about corruption—and electoral accountability.

Finally, electoral rules can also create mechanisms and incentives to increase political accountability. In a follow-up paper, Ferraz & Finan (2011) compare corruption practices for mayors audited at the same time but who differ in whether they are serving in a first term (eligible for re-election) or a second term (ineligible for re-election). Using the share of total federal resources transferred to municipalities associated with fraud in the public procurement of goods and services, diversion of funds, and overinvoicing of goods and services as a measure of corruption, they find that the share of stolen resources is, on average, 27% lower among mayors with re-election incentives than among mayors without re-election incentives. Although this result suggests that a two-term period is more effective than a one-term period as an anticorruption policy, it does not mean that politicians should be re-elected indefinitely. The absence of term limits by opening up the possibility of long-term entrenchment may encourage politicians to develop long-term relationships or policies that benefit them and their families. Term limits could also produce benefits if, in the absence of the pressure of being re-elected, politicians have better incentives to implement socially optimal policies with a long-term horizon.

Selection. Although the simple framework above suggests that the selection of who chooses to become a bureaucrat is potentially important, there is relatively little evidence on this point. Ferraz & Finan (2010) find that higher salaries attract better political candidates in Brazil, although the effects are relatively modest—a 20% increase in wages leads only to a 0.2 increase in the average years of schooling and a 0.05 increase in the

---

[9]Björkman & Svensson (2010) examine a community monitoring intervention in Uganda, in which local NGOs encouraged communities to be more involved in the state of health service provision. The intervention included meetings to discuss baseline information on the status of health service delivery relative to other providers and the government standard and encouraged community members to develop a plan identifying key problems and steps the providers should take to improve health service provision. The intervention increased the quality and quantity of primary health care provision; however, the design of the intervention suggests that the mechanisms could have included better information flows and monitoring.

number of terms of experience. Higher wages also improve the performance of a politician while in office. A 20% increase in wages leads to a 25% increase in the number of bills submitted, however. The authors do not, however, examine impacts on corruption per se. Evidence on the selection of politicians and the impact on corruption forms an important area for future work.

Selection based on the propensity to be corrupt may lead to multiple equilibria in corruption. In particular, in a corrupt equilibrium, the people who have the highest propensity to take advantage of corruption will disproportionately choose to become civil servants, which could make fighting corruption in the future more difficult—i.e., the same policies that effectively control corruption in a low-corruption country might not be enough to eliminate corruption in a high-corruption country, and in fact the same set of incentives might be consistent with both high- and low-corruption equilibria. Testing whether there are multiple equilibria in corruption—for reasons of selection or for other reasons—is an important area for future work.

Incentives. The framework outlined above is implicitly a model in which the only way the principal can observe what the bureaucrat is doing is through monitoring. For many government activities, however, there are direct indicators of agent behavior. This opens the possibility of tying incentives more closely to performance either through direct financial awards or more complex incentive schemes through promotions, assignments, and the like. However, performance indicators can be imperfect measures of the civil servants' corrupt behavior: We want citations issued only for those drivers who actually break the law, and not issued for those who do not; we want taxes collected when they are due, but we do not want overzealous tax collectors collecting from those who do not owe. Hence, in designing such incentive schemes, it is critical to deal with the so-called multitasking problem (Holmstrom & Milgrom 1991) and ensure that the true goals of the principal are achieved, not just the ones that are incentivized.

Evidence that directly links performance pay or other incentive schemes with corruption outcomes is largely lacking. Much of what we know comes from studies in the health and education sector, in which incentives are conditioned either on worker absenteeism or directly on health or education outcomes. The evidence on the effectiveness of incentive schemes that condition on worker absenteeism is mixed—in situations in which mechanisms to monitor performance are relatively tamper-free, performance-based pay can reduce absenteeism and, in the case of school teachers, improve test scores (Duflo et al. 2012). However, if service providers can collude with their monitors and/or tamper with monitoring tools, then such mechanisms may be undone in the medium run (Banerjee et al. 2008).

Similarly, evidence on the effect of performance pay for teachers is mixed. Muralidharan & Sundararaman (2011) conduct a randomized evaluation of a teacher incentive program in government-run rural primary schools in India. Teachers in treatment schools were eligible to receive a bonus payment based on the improvement of the students' average test scores. Some schools were assigned to a group incentive treatment, in which teachers were paid a group bonus based on improvements in the school-level average test score, whereas other schools were assigned to an individual incentive treatment in which the teacher was paid an individual bonus based on improvement in the average test scores of his students. They find significant improvements in student test scores. However, although group and individual incentive schools performed equally well in the first year of the program, average test scores were 0.10 standard deviations higher in schools where teachers were given individual

incentives relative to schools with group incentives by the end of second year of program. In contrast, Glewwe et al. (2010) show that performance pay for teachers in Kenya led to an improvement in outcomes only along the measures used to compute the formula that determines pay.

Outside education and health, there is little evidence on how incentives change the performance of bureaucrats. Kahn et al. (2001) use tax reform instituted by the Brazilian government in 1989 to study the effect of performance-based wages for tax collectors, in an economy with widespread tax evasions. The reform offered a bonus to tax officials based on group and individual performance in finding and collecting taxes from tax evaders. They find that the growth rate of fine collection exhibits a break in 1989 and estimate that fine collections per inspection are 75% higher on average than what they would have been in the absence of the program, with substantial heterogeneity across regions. The authors do their best to provide evidence that the surge in tax collections post-1989 resulted from the performance incentives provided by the tax reform, but because they do not have a control group, the evidence is suggestive but not conclusive.

Taken together, the tentative conclusion of this evidence is that there is room for incentives to succeed, but caution must be taken to design the incentives well and prevent them from being undermined. It is striking that few studies directly examine the impact of improved incentives on corruption outcomes.

3.1.2. The market for bribes: changing the structure of the bureaucracy to harness the forces of competition.
The previous section focuses on how a principal—i.e., the government—can best monitor its agents—civil servants. In other settings, however, strategic interactions between corrupt agents themselves become important, and depending on how the market is structured, these kinds of strategic interactions can either raise or lower the bribe amounts.

If a person needs to bribe multiple corrupt officials to perform a given task, Shleifer & Vishny (1993) argue that the "double-marginalization" problem can arise. Specifically, if each agent does not fully internalize the effect of his bribes on other agents' bribe revenues, the total amount of bribes one would need to pay could be higher than if agents had acted independently.

Olken & Barron (2009) use data on the bribes truck drivers pay to empirically test the idea that market forces partially determine the level of corruption and specifically to test for this type of double-marginalization. They exploit the fact that, during the period studied, the number of checkpoints along one of the roads was reduced in accordance with a peace agreement signed earlier in the year. They use this change in market structure to estimate the elasticity of the average bribe paid with respect to the expected number of checkpoints. They show that the average price paid at checkpoints increases when the number of checkpoints declines, consistent with the double-marginalization idea. These findings highlight the need to consider strategic interactions between corrupt agents themselves, in addition to interactions between principals and agents, in designing effective anticorruption policy.

An implication of this view is that a policy reform that moves from having a large number of independent agents to a single agent may reduce corruption and increase economic efficiency. Bruhn (2008) uses the sequential implementation of a reform that simplified business entry regulations across municipalities in Mexico to estimate the economic effects of such reforms more convincingly than in cross-country data. Although the paper does not look at the effects of the reform on corruption directly, the results show that simplified regulation

improved efficiency. She finds that the reform increased the number of registered businesses by 5%, which was accounted for by former wage earners opening businesses. Wage employment also increased by 2.2% as a result of the reform, whereas competition from new entrants decreased the income of incumbent businesses by 3%. It is, however, not possible to distinguish between less potential for corruption and increased convenience as the mechanism underlying efficiency gains.

The flip side of strategic interactions between bureaucrats is that if bureaucrats are competing against one another, this could reduce the bribes paid and lead to lower bribes and more output. In a recent study, Burgess et al. (2011) explore this issue in the context of deforestation in Indonesia. In particular, they explore a setting in which local district forestry officials can allow logging beyond the legal logging quota in exchange for bribes. The study shows that as the number of political jurisdictions increases, so that there are more bureaucracies with the potential to facilitate illegal logging in a province, logging rates increase and prices for wood fall, consistent with a model of Cournot competition between bureaucrats.

Despite the potential for competition between bureaucrats to reduce bribes, other than Burgess et al.'s (2011) study we know of no other evidence that examines how competition between bureaucrats works in practice. In Burgess et al.'s study, competition occurs only through the product market: Each district chooses how much wood to extract, and market forces—a common demand curve—determine how much they receive in rents. In many other settings, however, individual agents would be able to choose which bureaucrat to work with to obtain a service, and the bureaucrats might compete on price. This type of Bertrand competition could result in even larger impacts of competition than the type of Cournot-style competition studied by Burgess et al. (2011). We regard further studies examining competition between agents as a first-order question for future work.

It is important to note that competition leading to lower bribes is not necessarily socially optimal. In particular, it depends on what the government is trying to accomplish and whether the bribes are on top of, or instead of, official government fees. For example, in the case of deforestation studied by Burgess et al. (2011), bribes were to allow more logging than the government had deemed optimal (for example, for reasons of watershed protection or biodiversity protection). Competition meant lower bribes and greater quantities, which in this context meant more illegal logging, and hence greater social losses, than had there been less competition. Conversely, in the case of the road checkpoints studied by Olken & Barron (2009), traveling the road should have been free, so lower bribes would have meant greater road travel and greater efficiency. Understanding the welfare implications of these types of strategic interactions depends therefore on whether higher or lower bribes would increase or decrease social efficiency, and we do not yet know of an empirical example demonstrating how competition between bureaucrats could lead to greater social efficiency.

3.2. Transparency

One of the key themes of the international anticorruption movement is the role of transparency—so much so that the largest worldwide anticorruption nongovernmental organization (NGO) is called Transparency International. But does transparency matter?

The basic idea about transparency is that by enabling information about government actions, citizens can better monitor government officials and enforce greater electoral

accountability. However, the effect of making information about politicians publicly available is a priori unclear. Although disclosure of information can increase political accountability, it can also undermine politicians' privacy and thus potentially worsen the pool of entrants.

Several pieces of evidence suggest a relationship between access to information about politicians' performance and both the political accountability and the quality of government. In a cross-sectional, cross-country study, Djankov et al. (2010) study the relationship between disclosure rules for information about parliament members and a number of measures of government quality and corruption. Their main conclusion is that public disclosure, but not internal disclosure to parliament, is associated with lower perceived corruption and better government. They further find that information about politicians' assets, liabilities, income sources, and potential conflicts, as opposed to simply income and wealth levels, is more consistently associated with better government. Because the study is cross-sectional in nature, they cannot rule out reverse causality (i.e., higher-quality governments adopt better transparency laws).

In a more micro example, Banerjee et al. (2012) study how public disclosures about politicians' performance and qualifications can influence electoral accountability in settings characterized by weak institutions and a less educated population by conducting a randomized experiment in Delhi, India. Using the Indian Right to Information Act and candidates' affidavits, they created report cards for 10 assembly jurisdictions during the run-up to the 2008 election in Delhi. They then randomly provided slum dwellers with pamphlets and free newspapers containing information on candidate qualifications and legislator performance. The information increased voter turnout by 3.5 percentage points and reduced the incidence of vote buying by 19 percentage points. The information campaign seems to increase the quality of government: The vote share of the best-performing incumbent increased by 7 percentage points in the treatment group relative to the controls.

A related example, discussed in the monitoring section (Section 2.1.6), is Ferraz & Finan's (2008) study in Brazil on how public dissemination of corruption scandals in local governments had a negative effect on incumbents' electoral performance. Importantly, they find more pronounced impacts in areas where local radio stations were present to broadcast the results of the audit reports. The probability of re-election for an incumbent who committed two corruption violations in municipalities with pre-election audits was 7 percentage points lower than one who had zero violations and 11 percentage points lower if radio stations were present in the municipality. One interpretation for the larger effect in municipalities where local radios were present to divulge the information is that radios provide an efficient way of transmitting information about local politics.

A second way that transparency may matter—and the way that many suggest it does—is by providing citizens with information on what they are entitled to. Reinikka & Svensson (2005) study how an information campaign to monitor local officials can reduce corruption and increase educational outputs. They exploit a newspaper campaign in Uganda aimed at reducing the capture of public funds by providing students' parents with information to monitor local officials' use of an educational grant. Their empirical strategy uses distance to the nearest newspaper outlet as an instrument of school exposure to information and finds that an increase in information resulted in an increase in spending reaching the schools and ultimately an increase in school enrollment and student learning. An important caveat is that the distance to newspaper outlets may be nonrandomly assigned, and it may also have other, direct impacts on educational performance.

A third way in which transparency could matter is by allowing citizens to signal interest in a particular outcome. Peisakhin & Pinto (2010) examine this by conducting a randomized experiment to test whether freedom-of-information laws can improve access to basic public goods that are otherwise attainable only through bribery. The experiment randomly assigned individual applicants in India to one of three mechanisms used when requesting public benefits and then tested the effect of these mechanisms on the time that elapsed before the applicant received the benefit. In the first treatment group, applicants submitted an information request under the Right to Information Act shortly after their applications. The second group of applicants presented a letter of support from a local NGO with their application. Finally, the third group of applicants paid a bribe to a local to obtain the benefits. According to the results, 94% of those who paid bribes or sent an information request received benefits over the course of one year, as opposed to 21% in the NGO and control groups. Individuals in the group that paid bribes received benefits in a median of 82 days, 38 days less than those in the groups that filed an information request. The groups that neither paid a bribe nor requested information obtained benefits only after 343 days. The results suggest that requesting information under the freedom-of-information law is a reasonable, although imperfect, substitute for bribing an official. In a follow-up study, Peisakhin (2012) estimates the effect of the freedom-of-information law in the process of voter registration and here finds that the information law is an effective, free, and legal substitute to bribery for middle-class applicants.

3.3. Technology and Communications

Technological innovations can help make tools available that are hard for humans to tamper with and that enhance communication. For many (although not all) corrupt activities, the corrupt agent needs to somehow evade the rules or procedures that the official government bureaucracy has set up. Technology can help address this problem by ensuring mechanically that certain procedures are followed (Duflo et al. 2012).

Technology can also have a substantial impact on corruption by facilitating communication, which can enable better monitoring. Yang (2008) explores how hiring foreign inspectors to verify the tariff classification and the value of shipments before they leave their origin country impacts import duty collections. The key mechanism to reduce customs fraud is the transmission of information from the foreign firm at the origin port to the client government. The flow of information could not only improve the monitoring ability, but also reduce the bargaining power of corrupt customs officials, which can reduce bribe payments and custom clearance time. The results from the study provide evidence in favor of this hypothesis, showing that preshipment import inspection programs increased import duty collection by 15–30 percentage points in the first five years after implementation.

Relatedly, in many countries technology has played an important role in the design and administration of the tax system. One key idea of tax enforcement is double-reporting, in which the tax department compares two independent reports about tax performance and investigates discrepancies. In the developed world, Kleven et al. (2010) analyze a randomized tax-enforcement experiment in Denmark and find that the tax evasion rate is very small (0.3%) for income subject to double-reporting and much higher (37%) for self-reported income. Relatedly, in Chile, Pomeranz (2010) finds that audits have a much larger impact on the part of the VAT chain where there is not double-reporting, suggesting that for the rest of the VAT chain, double-reporting played an important role in encouraging

truth telling. Technology plays a key role here: In a manual system, actually doing the matching from all the double-reported information would be challenging, but once the system is automated it is much easier. Given the large number of countries in the process of modernizing and computerizing their tax infrastructure, it should be possible to study the impacts of this type of technology in the context of poorer countries where tax evasion is usually higher.

A final area in which technology holds promise is in procurement. Throughout the world, including in the developing world, governments have been moving to online procurement systems. These systems can potentially reduce corruption by increasing access to information (undermining bidding rings) and by making the procurement system more transparent. Lewis-Faupel et al. (2011) examine the impact of electronic procurement for road projects in India, taking advantage of a staggered rollout across Indian states. They find that electronic procurement leads to higher-quality roads, as measured by independent central government audits, although not to lower costs. The evidence suggests that the quality improvement comes from higher-quality contractors being more likely to win contracts.

4. SOME CAVEATS: ADAPTATION IN THE SHORT AND MEDIUM RUN

Much of the evidence discussed above shows the short-run effects of anticorruption policies and programs. But there is ample evidence to believe that the long-run impacts could be quite different. For example, it could take corrupt officials time to learn how to manipulate a new system, so the long-run effects of an anticorruption policy could be smaller than the short-run effects. Alternatively, it could take time for a new group of civil servants to select into the system, so an anticorruption policy could be more effective over time if it encourages more low-corruption types to select into the civil service. Or officials might simply substitute from one form of corruption to another.

In the short run, there are several examples of substitution from one type of corruption to another. In Olken's (2007) study, an increase in auditing of road expenditures led to decreased missing expenditures from the project, but to more family members of project officials being hired to build the roads. In Niehaus & Sukhtankar's (2010) study, conversely, an increase in the wages of daily wage jobs (and hence in the ability to steal from those workers) led to a reduction in theft of piece-rate jobs. Burgess et al. (2011) find that when a district's oil and gas revenues increase, providing an alternate source of rent extraction for local district officials, illegal logging falls. In all these cases, it appears that corrupt officials have different avenues of corruption available to them and substitute among them as the returns to one form of corruption get easier or harder. If an anticorruption policy clamps down differentially on certain types of corruption (as almost all do), one needs to take care that corrupt officials do not substitute to other forms of corruption with more severe efficiency costs.

There are also several examples that suggest that the long-run effect of anticorruption policies may be smaller than the short-run effect as officials adapt. One of the examples mentioned above is Banerjee et al. (2008). In this study, an incentive program on nurse attendance in India was found effective only during the first six months of the intervention, when the program was correctly in place. Later, however, the system was undermined by the local health administration, who took advantage of a loophole in the program design and began allowing many more unexcused absences. By 18 months after the program had started, the program was no longer able to improve nurse attendance.

In the case of Brazil, Ramalho (2007) uses the 1992 impeachment of President Fernando Collor to evaluate the impact and persistence of corruption on the market value of politically connected companies. The results suggest that the market perceived the decrease in the president's probability of staying in power as affecting the value of politically connected companies, but only temporarily. According to the results, family-connected companies had on average daily abnormal returns that were 2–9 percentage points lower during bad event days, with the effect reversing completely within one year. One interpretation is that over the course of the year, these previously politically connected firms were able to form new connections.

In Colombia, Camacho & Conover (2009) study the manipulation of the poverty index score as an eligibility requirement to gain access to social programs. In particular, people were eligible for the program if their poverty index was below a given threshold. In the first year of the program, there was no discontinuity in poverty index scores at the threshold, but over the subsequent years, as the formula became better known, more and more manipulation began to take place, resulting in a larger number of people with scores just below the eligibility threshold. The results suggest that in total 3 million people had their score changed, which accounts for approximately 40% of the beneficiaries.

A final example comes from Burgess et al. (2011). As noted above, when a district's oil and gas revenues increase, providing an alternate source of rent extraction for local district officials, illegal logging falls. However, within three years, the effect reverses and illegal logging returns to almost its previous level. Burgess et al. provide suggestive evidence that the mechanism is a change in the political equilibrium—the higher oil and gas rents change the nature of the governing coalition toward a type of coalition associated with higher rent extraction. This new political coalition presumably extracts rents not just from oil and gas, but also from the forest sector.

5. CONCLUDING THOUGHTS

Recent years have seen a dramatic rise in microempirical research on corruption in developing countries. A string of papers has shown how basic economic concepts can be applied to corruption. Corrupt officials respond to incentives and the threat of punishment, even in corrupt environments. Strategic interactions between corrupt officials affect the level of corruption—bidding down bribes if they compete against one another, and increasing bribes if multiple bribes are required and officials cannot coordinate with one another. Recent work has also shown that corrupt officials are resilient: Over time, they adapt to changes in their environments, in some cases offsetting anticorruption policies with new avenues for seeking out rents.

Although these examples highlight cases in which adaptation took place over time, it is also possible that in some situations the long-term impacts of anticorruption policies exceed the short-run effects. One area where this is likely is the case of transparency reforms: A higher likelihood of information disclosures may both incentivize politicians to perform better and improve the incentives for high-talent individuals to enter politics.

It is also important to note that there is relatively little research on many of the main anticorruption policy initiatives we observe in the world today. When aiming to fight corruption domestically, many countries set up high-profile, independent anticorruption agencies with prosecutorial powers. From the international perspective, there have been major efforts at promoting transparency (such as the Extractive Industries Transparency Initiative), and at limiting the ability of foreign companies to pay bribes (such as the OECD

Anti-Bribery Convention and the US Foreign Corrupt Practices Act). Despite the recent advances in research on corruption, we know relatively little about whether these types of domestic and international policy initiatives are successful, and how potentially corrupt agents respond to these types of policies.

Taken together, the existing and ongoing research on corruption in developing countries suggests a large and promising research agenda. Researchers have identified several innovative ways of measuring corruption, and economic theory offers us significant guidance on how to design anticorruption policies. At the same time, the ability of individuals to outguess those who seek to regulate them suggests an important need to collect data on both the short- and long-run impacts of many different anticorruption policies in many different contexts.

DISCLOSURE STATEMENT

ACKNOWLEDGMENTS

LITERATURE CITED

Bandiera O, Prat A, Valletti T. 2009. Active and passive waste in government spending: evidence from a policy experiment. *Am. Econ. Rev.* 99:1278–308

Banerjee A. 1997. A theory of misgovernance. *Q. J. Econ.* 112:1289–332

Banerjee A, Duflo E, Glennerster R. 2008. Putting a band-aid on a corpse: incentives for nurses in the Indian public health care system. *J. Eur. Econ. Assoc.* 6:487–500

Banerjee A, Hanna R, Mullainathan S. 2012. Corruption. In *Handbook of Organizational Economics*, ed. R Gibbons, J Roberts. Princeton, NJ: Princeton Univ. Press. In press

Banerjee AV, Pande R. 2009. *Parochial politics: ethnic preferences and politician corruption*. Work. Pap., Harvard Univ.

Becker GS, Stigler GJ. 1974. Law enforcement, malfeasance, and compensation of enforcers. *J. Legal Stud.* 3(1):1–18

Bertrand M, Djankov S, Hanna R, Mullainathan S. 2007. Obtaining a driver's license in India: an experimental approach to studying corruption. *Q. J. Econ.* 122:1639–76

Besley T, Pande R, Rao V. 2012. Just rewards? Local politics and public resource allocation in South India. *World Bank Econ. Rev.* 26:191–216

Björkman M, Svensson J. 2010. When is community-based monitoring effective? Evidence from a randomized experiment in primary health in Uganda. *J. Eur. Econ. Assoc.* 8:571–81

Brollo F, Nannicini T, Perotti R, Tabellini G. 2010. *The political resource curse*. NBER Work. Pap. 15705

Bruhn M. 2008. *License to sell: the effect of business registration reform on entrepreneurial activity in Mexico*. Policy Res. Work. Pap. 4538, World Bank

Burgess R, Hansen M, Olken B, Potapov P, Sieber S. 2011. *The political economy of deforestation in the tropics*. Unpublished manuscript, Mass. Inst. Technol.

Camacho A, Conover E. 2009. *Manipulation of social program eligibility: detection, explanations and consequences for empirical research*. Unpublished manuscript, Univ. Andes

Caselli F, Michaels G. 2009. Do oil windfalls improve living standards? Evidence from Brazil. *Discuss. Pap.*, Cent. Econ. Perform., London Sch. Econ.

Chaudhury N, Hammer J, Kremer M, Muralidharan K, Rogers FH. 2006. Missing in action: teacher and health worker absence in developing countries. *J. Econ. Perspect.* 20(1):91–116

Di Tella R, Schargrodsky E. 2004. Do police reduce crime? Estimates using the allocation of police forces after a terrorist attack. *Am. Econ. Rev.* 94:115–33

Djankov S, La Porta R, Lopez-de-Silanes F, Shleifer A. 2010. Disclosure by politicians. *Am. Econ. J. Appl. Econ.* 2(2):179–209

Duflo E, Hanna R, Ryan S. 2012. Incentives work: getting teachers to come to school. *Am. Econ. Rev.* In press

Faccio M. 2006. Politically connected firms. *Am. Econ. Rev.* 96:369–86

Ferraz C, Finan F. 2008. Exposing corrupt politicians: the effects of Brazil's publicly released audits on electoral outcomes. *Q. J. Econ.* 123:703–45

Ferraz C, Finan F. 2010. *Motivating politicians: the impacts of monetary incentives on quality and performance*. Unpublished manuscript, Univ. Calif., Berkeley

Ferraz C, Finan F. 2011. Electoral accountability and corruption: evidence from the audit reports of local governments. *Am. Econ. Rev.* 101:1274–311

Ferraz C, Finan F, Moreira DB. 2010. *Corrupting learning: evidence from missing federal education funds in Brazil*. Unpublished manuscript, Univ. Calif., Berkeley

Fisman D, Fisman R, Galef J, Khurana R. 2006. *Estimating the value of connections to Vice-President Cheney*. Unpublished manuscript, Columbia Bus. Sch., Columbia Univ.

Fisman R. 2001. Estimating the value of political connections. *Am. Econ. Rev.* 91:1095–102

Fisman R, Svensson J. 2007. Are corruption and taxation really harmful to growth? Firm level evidence. *J. Dev. Econ.* 83:63–75

Fisman R, Wei SJ. 2004. Tax rates and tax evasion: imports in China. *J. Polit. Econ.* 112:471–96

Glewwe P, Nauman I, Kremer M. 2010. Teacher incentives. *Am. Econ. J. Appl. Econ.* 2(3):205–27

Gordon R, Li W. 2009. Tax structures in developing countries: many puzzles and a possible explanation. *J. Public Econ.* 93:855–66

Gorodnichenko Y, Peter KS. 2007. Public sector pay and corruption: measuring bribery from micro data. *J. Public Econ.* 91:963–91

Holmstrom B, Milgrom P. 1991. Multitask principal-agent analyses: incentive contracts, asset ownership, and job design. *J. Law Econ. Organ.* 7:24–52

Hsieh CR, Moretti E. 2006. Did Iraq cheat the United Nations? Underpricing, bribes, and the Oil for Food Program. *Q. J. Econ.* 121:1211–48

Hunt J. 2007. How corruption hits people when they are down. *J. Dev. Econ.* 84:574–89

Kahn CM, Silva ECD, Ziliak JP. 2001. Performance-based wages in tax collection: the Brazilian tax collection reform and its effects. *Econ. J.* 111:188–205

Khwaja AI, Mian A. 2005. Do lenders favor politically connected firms? Rent provision in an emerging financial market. *Q. J. Econ.* 120:1371–411

Kleven HJ, Knudsen MB, Kreiner CT, Pedersen S, Saez E. 2010. *Unwilling or unable to cheat? Evidence from a randomized tax audit experiment in Denmark*. NBER Work. Pap. 15769

Lewis-Faupel S, Neggers Y, Olken B, Pande R. 2011. *Can electronic procurement improve infrastructure provision? Evidence from a large rural road program in India*. IGC Work. Pap., Int. Growth Cent., London Sch. Econ.

Malesky E, Samphantharak K. 2008. Predictable corruption and firm investment: evidence from a natural experiment and survey of Cambodian entrepreneurs. *Q. J. Polit. Sci.* 3:227–67

Mauro P. 1995. Corruption and growth. *Q. J. Econ.* 110:681–712

McMillan J, Zoido P. 2004. How to subvert democracy: Montesinos in Peru. *J. Econ. Perspect.* 18(4): 69–92

Mocan N. 2008. What determines corruption? International evidence from microdata. *Econ. Inq.* 46:493–510

Monteiro J, Ferraz C. 2010. *Does oil make leaders unaccountable? Evidence from Brazil's offshore oil boom*. Unpublished manuscript, PUC-Rio

Muralidharan K, Sundararaman V. 2011. Teacher performance pay: experimental evidence from India. *J. Polit. Econ.* 119:39–77

Niehaus P, Atanassova A, Bertrand M, Mullainathan S. 2012. Targeting with agents. *Am. Econ. J. Econ. Policy.* In press

Niehaus P, Sukhtankar S. 2010. *Corruption dynamics: the golden goose effect*. Work. Pap., Univ. Calif., San Diego

Olken BA. 2006. Corruption and the costs of redistribution: micro evidence from Indonesia. *J. Public Econ.* 90:853–70

Olken BA. 2007. Monitoring corruption: evidence from a field experiment in Indonesia. *J. Polit. Econ.* 115:200–49

Olken BA. 2009. Corruption perceptions vs. corruption reality. *J. Public Econ.* 93:950–64

Olken BA, Barron P. 2009. The simple economics of extortion: evidence from trucking in Aceh. *J. Polit. Econ.* 117:417–52

Olken BA, Pande R. 2011. Governance review paper: J-PAL governance initiative. *Rep.*, Abdul Latif Jameel Poverty Action Lab, Cambridge, MA

Peisakhin L. 2012. Transparency and corruption: evidence from India. *J. Law Econ.* In press

Peisakhin L, Pinto P. 2010. Is transparency an effective anti-corruption strategy? Evidence from a field experiment in India. *Regul. Gov.* 4(3):261–80

Pomeranz D. 2010. *No taxation without information: deterrence and self-enforcement in the value added tax*. Unpublished manuscript, Harvard Univ.

Ramalho R. 2007. The persistence of corruption: evidence from the 1992 presidential impeachment in Brazil. *Rep.*, World Bank, Washington, DC

Rauch JE, Evans PB. 2000. Bureaucratic structure and bureaucratic performance in less developed countries. *J. Public Econ.* 75:49–71

Reinikka R, Svensson J. 2004. Local capture: evidence from a central government transfer program in Uganda. *Q. J. Econ.* 119:679–706

Reinikka R, Svensson J. 2005. Fighting corruption to improve schooling: evidence from a newspaper campaign in Uganda. *J. Eur. Econ. Assoc.* 3:259–67

Sequeira S, Djankov S. 2010. *An empirical study of corruption in ports*. MPRA Pap. 21791, Univ. Munich

Shleifer A, Vishny R. 1993. Corruption. *Q. J. Econ.* 108:599–617

Svensson J. 2003. Who must pay bribes and how much? Evidence from a cross section of firms. *Q. J. Econ.* 118:207–30

Svensson J. 2005. Eight questions about corruption. *J. Econ. Perspect.* 19(5):19–42

Treisman D. 2000. The causes of corruption: a cross-national study. *J. Public Econ.* 76:399–457

Van Rijckeghem C, Weder B. 2001. Bureaucratic corruption and the rate of temptation: Do wages in the civil service affect corruption, and by how much? *J. Dev. Econ.* 65:307–31

Wei SJ. 2000. How taxing is corruption on international investors? *Rev. Econ. Stat.* 82(1):1–11

Yang D. 2008. Integrity for hire: an analysis of a widespread customs reform. *J. Law Econ.* 51(1):25–57

Zitzewitz E. 2012. Forensic economics. *J. Econ. Lit.* In press