European Economic Review 94 (2017) 1–22

Contents lists available at ScienceDirect

European Economic Review

journal homepage: www.elsevier.com/locate/euroecorev

I paid a bribe: An experiment on information sharing and extortionary corruption☆

Dmitry Ryvkin[a,*], Danila Serra[b], James Tremewan[c]

[a] *Department of Economics, Florida State University, 113 Collegiate Loop, Tallahassee, FL 32306-2180, United States*
[b] *Department of Economics, Southern Methodist University, 3300 Dyer Street, Suite 301C Umphrey Lee Center, Dallas TX 75275-0496, United States*
[c] *Department of Economics, University of Vienna, Oskar-Morgenstern-Platz 1, A-1090 Vienna, Austria*

ARTICLE INFO

Article history:
Received 16 May 2016
Accepted 7 February 2017
Available online 16 February 2017

JEL Classification:
D73
D49
C91

Keywords:
Information sharing
Extortionary corruption
Experiment
Crowdsourcing

ABSTRACT

Theoretical and empirical research on corruption has flourished in the last three decades; however, identifying successful anti-corruption policies remains a challenge. In this paper we ask whether bottom-up institutions that rely on voluntary and anonymous reports of bribe demands, such as the *I paid a bribe* website first launched in India in 2010, could act as effective anti-corruption tools, and, if this is the case, whether and how their effectiveness could be improved. We overcome measurement and identification problems by addressing our research questions in the laboratory. Our results show that the presence of a reporting platform like the *I paid a bribe* website may be insufficient to systematically lower bribery. A more effective platform is one where posts disclose specific information about the size of the bribes and the location of their requestors, i.e., a platform that could serve as a search engine for the least corrupt officials, especially if posting is restricted to service recipients. Our results also show that while citizens rarely post false information, lying by officials, when allowed to post on the platform, is widespread.

1. Introduction

Corruption is widespread around the world.[1] The most recent Global Corruption Barometer report by Transparency International[2] highlights that more than 25% of 114,000 survey respondents in 107 countries report having paid a bribe in the last 12 months when dealing with officials in at least one of eight public sectors, with the police, the judiciary and registration offices being the most corrupt. The percentage of bribe payers is well above 50% in several countries, including India, Liberia, Sierra Leone and Uganda.

☆ We thank two anonymous referees and seminar participants at Fordham University, the Higher School of Economics, the New Economic School, Southern Methodist University, the University of Texas at Dallas and the North-American Economic Science Association meeting for useful comments and suggestions. We also thank Philip Brookins for excellent assistance with the programming of the experiments.

* Corresponding author.
E-mail addresses: dryvkin@fsu.edu (D. Ryvkin), dserra@smu.edu (D. Serra), james.tremewan@univie.ac.at (J. Tremewan).

[1] For the most recent country-level corruption ranking, see Transparency International's 2015 Corruption Perception Index at http://www.transparency.org/cpi2015/results.

[2] http://www.transparency.org/gcb2013/in_detail.

http://dx.doi.org/10.1016/j.euroecorev.2017.02.003
0014-2921/

Evidence that corruption is harmful to society abounds.[3] It is, therefore, not surprising that studies of its causes and possible remedies have proliferated in recent years. Numerous methodologies have been employed to study corrupt behavior and identify effective anti-corruption policies; from theoretical models to cross-country comparisons to field and lab experiments.[4] The consensus is that corruption can be successfully reduced by employing monitoring and enforcement institutions that increase its expected costs by relying on high probabilities of punishment and severe sanctions. But how can this be achieved in societies characterized by systemic corruption where top-down monitoring and rule enforcement can be easily bypassed through the payment of bribes, and where corrupt exchanges have long been embedded in the prevailing social norms? In these societies, relying on institutions that attempt to curb corruption from the bottom-up, for instance by encouraging anonymous reports from the recipients of corrupt goods and services, may be the only viable solution. It is often argued (see, e.g., The World Bank, 2004) that bottom-up initiatives, although unable to affect public officials' incentives through the threat of monetary punishment, might work because service recipients may be more willing than top-down inspectors to monitor corruption in order to avoid the costs that they would otherwise incur from it, and because public officials may want to avoid the nonmonetary costs generated by social disapproval from members of their community.

Starting from the launch of the *I paid a bribe* website[5] in India in 2010, anti-corruption initiatives promoting online reports of corrupt exchanges have appeared in a number of developing countries, including Kenya, Indonesia and Pakistan. The *I Paid a Bribe* website is self-described as focusing on "crowdsourced reports of corruption." The website gives citizens the opportunity to anonymously report their bribery experiences. After its launch in India in 2010, the website is now operative in a total of 14 countries around the world. Despite the increasing popularity of these corruption reporting platforms,[6] their effectiveness in the fight against corruption is unknown. The *I paid a bribe* website, in particular, was designed with the aim of increasing observability and awareness of corruption and was not intended to serve as an actual anti-corruption tool. Its implementation at the national level, together with the difficulty of gathering objective measures of corruption before and after the appearance of the website, makes it impossible to reliably measure its *causal* impact on corruption. Scientific evaluations of similar bottom-up anti-corruption initiatives relying on citizens' reports are also missing. The only related study is that of Olken (2007), who assessed the effectiveness of citizen monitoring on corruption in the construction of public roads in Indonesia. However, in this study, monitoring occurred through participation in village meetings and through previously collected anonymous comment notes to be read at such meetings.[7] Laboratory experiments on the impact of bottom-up monitoring on corruption are equally scarce. While there are studies (Banuri and Eckel, 2015; Cameron et al., 2009; Serra, 2012) on the impact of top-down symmetric (Abbink et al., 2002) and asymmetric (Abbink et al., 2014) penalties, as well as bottom-up monetary sanctions, i.e., sanctions imposed by the victims of corruption on the perpetrators, on bribery, only Salmon and Serra (2016) examine the impact of social nonmonetary judgment on corruption. However, in Salmon and Serra (2016) the messages of social disapproval are seen in private by the corrupt individual who received them and are not displayed to the public as it is the case in the reporting platforms that we intend to evaluate in this paper.

None of the existing studies, to the best of our knowledge, attempt to investigate the effectiveness of bottom-up anti-corruption initiatives that rely on citizens' corruption reports in the same vein as the *I paid a bribe* website. This paper fills the gap in the literature by employing a laboratory experiment to assess whether bottom-up institutions that rely on voluntary and anonymous reports of bribe demands could act as an effective anti-corruption tool, and, if this is the case, how their impact could be maximized. We assess the effectiveness of a reporting platform that simulates the *I paid a bribe* website, and we then propose a number of changes to such a platform and test whether the resulting reporting systems could be more effective in lowering bribe demands and bribe payments.

The use of a laboratory experiment presents an advantage over field studies as it makes it possible to clearly observe the extent to which corrupt behavior responds to different bottom-up reporting systems. It could be argued that decisions made by student subjects within the controlled environment of the experimental lab are unlikely to reflect the decisions that the same students or ordinary citizens would make in outside-the-lab situations. While we acknowledge that the levels of corruption are likely to be different when using different subject pools and when moving from the lab to the field, like most experimental economists, we are not interested in levels; we are interested in comparative statics regarding the effects of certain incentive systems and contextual situations, and in the mechanisms underlying observed behaviors. There is substantial evidence of the external validity of qualitative results generated by lab experiments (Camerer, 2011; Kessler and Vesterlund, 2011). With respect to corruption, Armantier and Boly (2013) have shown that corruption "can be studied

[3] Corruption slows economic growth and development (Mauro, 1995; Méon and Sekkat, 2005), exacerbates inequality and poverty (Gupta et al., 1998), distorts the allocation of public spending away from education and health (Reinikka and Svensson, 2004; Tanzi and Davoodi, 1998), and impairs the provision of social services or programs to those most in need (Olken, 2006).

[4] For a review of important theoretical issues, see Bardhan (1997) and Ryvkin and Serra (2012). For examples of studies using cross-country data, see Treisman (2000) and Serra (2006). For examples of studies using household-level data, see Hunt (2007). For examples of studies using direct observation in the field, see Sequeira and Djankov (2013). For a review of experimental studies of corruption with clear policy implications, see Serra and Wantchekon (2012). For a critical assessment of the different methodologies employed for the empirical study of corruption, see Sequeira (2012).

[5] See http://www.ipaidabribe.com.

[6] For a recent article in the popular media advocating the use of corruption reporting platforms see, for example, http://www.theguardian.com/global-development-professionals-network/2015/may/19/crowdsourcing-anti-corruption-bribery-kenya-india.

[7] This form of bottom-up monitoring proved unsuccessful in reducing corruption. Partly related are a number of recent evaluations of the impact of community participation in education (Banerjee et al., 2010; Chaudhury et al., 2006; Duflo et al., 2014; Pradhan et al., 2014), and health (Björkman and Svensson, 2010).

in the lab" by comparing students' behavior and responses to incentives in a bribery experiment conducted in Canada, and non-student behavior and responses to incentives in the same laboratory experiment in Burkina Faso and in a field experiment, also in Burkina Faso. The results show that individuals' propensities to engage in bribery were virtually identical in the lab and in the field in Burkina Faso, and the treatment effects do not statistically differ across the three environments when controlling for individual characteristics.[8] The most important argument in favor of our methodology is that it makes it possible to investigate a question – i.e., whether a platform that allows citizens to report bribe demands could reduce corruption – that, due to the secretive and illegal nature of corruption, is impossible to address in the field. Moreover, the lab setting allows us to assess the extent to which false reports, by citizens or even officials, constitute a problem. Given the anonymous nature of the reports, this would be impossible to achieve in the field.

In this paper, like the *I paid a bribe* websites, we focus on the kind of corruption "that confronts ordinary citizens in their daily lives when they're not able to avail of services they are legitimately entitled to from the government – getting a driver's license, a birth certificate, registering a purchase of property."[9] This kind of corruption is often referred to as *extortionary* or *coercive* corruption, as opposed to *collusive* corruption, which takes place when a bribe is exchanged for the provision of an illegal good or service, for instance for the provision of a building permit to an unqualified firm.[10] In our experimental setting, citizens need to obtain a license and public officials can ask for bribes on top of the official licensing fee. Citizens can obtain the license from any of the available offices displayed on a map; however, every time they visit a new office they incur a cost. Naturally, this set-up applies to environments where more than one provider is available for the provision of a good or a service. Even in countries where service provision is primarily characterized by supply-side monopoly, there exist citizen-official transactions where citizens have a choice between a number of providers. This is typically the case for transactions that involve healthcare providers, traffic police and school choice.[11]

We compare our baseline setting to an *I paid a bribe* treatment (BB_IPB) where citizens can anonymously post on a public bulletin board (BB) the bribe(s) that they were asked to pay for the license. This comparison allows us to assess whether the *I Paid a bribe* website, in its current form, can serve as an effective anti-corruption tool. We then look into ways in which the *I Paid a Bribe* website could be improved without relying on top-down monitoring. In particular, we investigate the impact that an online reporting platform would have on bribe demands if it could also be used as a search engine for low-bribe-demanding officials.[12] To this end, we designed a treatment (*BB_CIT*) where citizens can leave posts not only about the size of the bribe demanded but also about the specific location of the office/official that demanded the bribe, hence effectively reducing each other's need to physically search for the least corrupt official. Ryvkin and Serra (2015), employing the same baseline corruption setting as in the current study, found that decreasing citizens' search costs lowers bribe demands. Similar effects have been found in the context of standard markets with respect to the relationship between search costs and prices (Cason and Friedman, 2003; Davis and Holt, 1996; Morgan et al., 2006). This led us to believe that if online reporting websites could be used to decrease search costs – or at least are perceived as such by public officials – they may constitute a relatively inexpensive way to significantly reduce corruption. The comparison between the BB_IPB and the BB_CIT treatments makes it possible to assess the potential benefits of redesigning the existing websites in such a way that citizens could effectively use them as information sharing mechanisms.

Finally, we are able to investigate whether false reports are a reason for concern. This is an important question that applies to any crowdsourcing reporting mechanism and is especially difficult to answer using field data. In our experimental setting we can easily check the truthfulness of the messages posted by citizens. Moreover, in the BB_IPB treatment and in a third treatment, BB_ALL, we also allow officials to post (possibly false) messages on the board. This way, we are able to assess the extent to which online reporting platforms are prone to information bias due to reports posted by officials (or by third parties acting in the interest of the officials).

Since the bribes demanded by public officials can be viewed as "prices" charged for the provision of a good or service, our study is also related to the literature on the effectiveness of online feedback and information sharing platforms in

[8] Additionally, while proving the external validity of our design is virtually impossible, some conclusions can be drawn from the study of Barr and Serra (2010), which relies on a bribery game similar to the one used in this paper and on a sample of students coming from over 40 countries characterized by markedly different levels of corruption. Barr and Serra (2010) show that behavior of the undergraduate students in the game could be predicted by the level of corruption in the students' home countries, as proxied by Transparency International's Corruption Perception Index. The fact that corruption in the stylized setting of the lab experiment correlates with corruption in the participants' home countries can be interpreted as an indication that the setting reproduced in the lab is indeed related to corruption decision-making outside the lab.

[9] This is a quote from the *I paid a bribe* website, see http://www.ipaidabribe.com.

[10] In the case of collusive corruption, a corrupt transaction may impose negative externalities on the society; think, for example, of a poorly constructed bridge that passes inspections and then breaks down. In the case of extortionary corruption, citizens have to pay bribes for services they are entitled to receive; therefore, the transaction does not generate negative externalities to others.

[11] The extent to which other services and sectors are characterized by multiple providers located in different geographical areas depends on the country and the sector under consideration. In sectors where officials have monopoly power over service delivery, the effectiveness of the citizen reporting platforms that we assess in this paper is conditional on the implementation of changes in the structure of the bureaucracy that introduce competition among providers, as studied in Ryvkin and Serra (2015, 2016).

[12] The *I paid a bribe* website was not designed to be used as a search engine. Consequently, when reporting a case of corruption, the website only requires information about the sector of the government where corruption occurred, the bribe demanded and possibly paid, and the city where the transaction occurred. No information about the *specific office* where corruption occurred is asked, and although citizens could voluntarily provide this information, the online system is such that this information would simply remain in individual stories and could not be used to search for honest officials when having to apply for a given service.

standard markets.[13] The existing literature shows that customers' ratings significantly affect sales (Anderson and Magruder, 2012; Cabral and Hortacsu, 2010; Chevalier and Mayzlin, 2006; Ye et al., 2009) and prices (Cabral and Hortacsu, 2010; Houser and Wooders, 2006; Resnick et al., 2006). A number of studies have also investigated strategic feedback in two-way rating systems (Bolton et al., 2013; Masclet and Pénard, 2012; Resnick et al., 2006) and the problem of fake reviews (Anderson and Simester, 2013; Luca and Zervas, 2013; Mayzlin et al., 2012). Recently, there have also been attempts to use experimental data (Lafky, 2014; Rockenbach and Sadrieh, 2012) to gain a better understanding of the motivations behind the decision to provide feedback about a seller. Contrary to the existing studies, our reporting platform has to do with prices (i.e., bribes). This is a unique feature of government service provision as compared to standard markets, since in the former case "prices," i.e., bribes, cannot be openly advertised, hence the added value of information sharing mechanisms. An important consequence of the fact that in our setting reporting is about prices and not quality is that the identification of posts containing false information is straightforward.

Our results suggest that the presence of a reporting platform such as the *I paid a bribe* website does not systematically reduce bribe demands and bribe payments. However, if the website could be modified so that the posts contained information not only on the size of the bribes but also on the location of the offices where such bribes where demanded, harassment bribes would be significantly lower. Our results show that such a website is likely to succeed even when the utilization rate by citizens is relatively low – about 40% in the experiment – as compared to that by officials – about 70%. Restricting the use of reporting systems to service recipients further reduces corruption. This is because when officials are allowed to leave messages on the reporting platforms, they tend to post multiple false messages, compromising in this way the credibility and efficacy of the reporting system. On the other hand, we see very little lying from citizens. As a consequence, the most effective reporting platform is one where only citizens are allowed to post and where the information posted is specific about both the size of the bribe demands and the location of the offices/officials where such demands originated. We believe that finding ways to implement such a platform in the field should be made a priority in the fight against corruption.

The rest of the paper is structured as follows. Section 2 describes our extortionary bribery experiment, the treatments and their implementation. Section 3 presents a theoretical framework that we employ to generate our predictions. Section 4 follows with the experimental results and Section 5 concludes.

2. The extortionary bribery experiment

We employ the extortionary bribery game first introduced by Ryvkin and Serra (2015), in which we use corruption-loaded language. In the experiment, subjects are randomly assigned either the role of public official or the role of private citizen and keep that role for the duration of the experiment. There are 7 citizens and 7 officials in each session. Each public official receives a lump-sum wage of 130 experimental currency units (ECU) and is in charge of an office that provides licenses to private citizens. The official license fee is 20 ECU; however, as in real life, this fee is not pocketed by the official. At the beginning of each round of the experiment, public officials simultaneously and independently decide whether to demand a bribe for the provision of the license on top of the official fee, and the size of the bribe, if any. Those officials who decide to demand a bribe can demand any amount between 1 and 50 ECU, and cannot change their bribe demands during the course of a round. They have the chance to modify their decision to demand a bribe and the size of the bribe, if any, only at the beginning of a new round.

Subjects play for a total of 10 rounds, although they are not told this in advance. Instead, at the end of each round except last, subjects are presented with an intermediate computer screen that informs them that they will be playing another round of the game. On the same screen, participants are also told that in the following round officials are randomly re-assigned to different offices.[14] We also change the labeling of offices from one round to the next. This implies that, in the experiment, citizens cannot associate location on the map to a particular official. In other words, since we are interested in studying corrupt transactions that are likely to be one-shot, like the ones targeted by the *I paid a bribe* websites, we do not allow citizens and officials to build long-term relationships.

In each round, citizens receive an initial endowment of 80 ECU and when they acquire the license they gain an additional 70 ECU. At the beginning, each citizen is randomly matched with an office so that 7 citizen-official pairs are formed. Each citizen finds out the bribe demanded by the corresponding official, if any. The citizen can either get the license by paying the license fee and the requested bribe, or visit another office by paying a fixed search cost of 5 ECU. Each citizen has access to a map showing all available offices; for the offices the citizen visited, the map shows what bribes are demanded there. Citizens can always go back, without paying the search cost, and get the license from any of the previously visited offices by paying the corresponding bribe; they have to get the license eventually. Officials do not know the size of the bribes

[13] Much online information on prices is gathered by websites that act as automatic search engines. However, here we are interested in information that is volunteered by participants in the transaction. Therefore, although our study is about the sharing of information about prices (i.e., bribes), studies on the effectiveness of sharing information about product quality (such as product ratings, recommendations and reviews) are the most relevant to our research endeavor.

[14] In the message displayed to the subjects we refer to a round as a "sequence." The message reads: "The next sequence is about to begin. In this sequence, you will be in exactly the same environment as in the previous sequence. You will retain the role of either Private Citizen or Public Official. However, note that all the Offices have changed and the Public Officials have been randomly re-assigned to different Offices."

Table 1
Summary of experimental sessions and treatments.

Treatment	Sessions	Subjects	Allowed to read	Allowed to post	Information posted
No_BB	6	84			
BB_CIT	6	84	All	Citizens	Specific
BB_ALL	6	84	All	All	Specific
BB_IPB	6	84	All	All	Not specific
Total	24	336			

demanded by other officials; however, as citizens search through the offices, each official can see how many citizens visited his or her office, how many decided to get the license there and how many decided to leave.

In the absence of corruption both public officials and citizens earn 130 ECU. The official's payoff at the end of each round is 130 ECU plus all the bribes paid by the citizens who decided to get the license at the official's office. The citizen's payoff at the end of each round is 150 ECU minus the 20 ECU cost of the license, the bribe she ends up paying and the total search costs, i.e., 5 ECU multiplied by the total number of searches.

In addition to the baseline treatment described above, we conduct three treatments in which we introduce a bulletin board (BB) that provides subjects an opportunity to post publicly visible structured messages reporting the size of the bribes demanded for the provision of the license. In our BB treatments, we manipulate the richness of the information displayed on the board and the set of the subjects that are allowed to post. In particular, in the BB_IPB treatment, the posts contain minimal information, as they only report the size of the bribes and not the specific office where each bribe was demanded. Moreover, in an attempt to replicate the *I paid a bribe* websites, we allow both citizens and officials to post. However, the way messages are pre-structured makes it clear that the BB is meant to be used by citizens. The message subjects can post reads: "I visited an office and was asked to pay a bribe of [*blank*] ECU." Participants have to fill in the blank and click "Enter Message" for the message to appear on the board.

In a second BB treatment, *BB_ALL*, both citizens and officials can post on the board, but we allow for more specific information to be posted. In particular, the messages now report both the size of the bribe and the office where the bribe was demanded. The message now reads: "I visited office [*blank*] and was asked to pay a bribe of [*blank*] ECU." Finally, in a third treatment, *BB_CIT*, we restrict the active use of the board to the citizens while keeping the information displayed in the posts as specific as in *BB_ALL*. In this treatment, while officials cannot post on the board, they can still view the information displayed on it.[15]

In all the BB treatments, we allow subjects to post false information on the board. Since we did not want to suggest or induce lying behavior through our experimental instructions, we did not mention the possibility of lying when explaining the rules of the game. Instead, we let subjects discover this possibility by themselves. As in the actual *I paid a bribe* website, subjects posting on the board and wishing to lie would soon realize that they could do that, since any message, no matter its content, would go through.[16]

Finally, in all BB treatments we wanted to avoid the situation where all citizens, after discovering the bribe requested by their randomly assigned office, would wait for the information about other offices to appear on the bulletin board before deciding whether to get the license in their current office or go somewhere else.[17] To this end, we allow citizens to observe the information on the bulletin board only after they have made their first pay/search decision. After that first pay or search decision, they can access the board at any time.

The experiment results in four treatments: baseline (*NO_BB*), *BB_CIT*, *BB_ALL* and *BB_IPB*. A total of 336 subjects (52.7% female) participated in 24 sessions of the experiment, as summarized in Table 1. Each experimental subject participated in only one session and, hence, one treatment.

Instructions were read aloud, with a printed copy distributed to subjects (sample instructions are provided in the Appendix). Before engaging in the corruption experiment each subject was involved in a task aimed at measuring risk preferences. Following the method first introduced by Holt and Laury (2002), we invited subjects to choose between two lotteries, $A = (\$1.60, \$2.00; p, 1-p)$ and $B = (\$0.10, \$3.85; p, 1-p)$, with probability p changing from 0 to 0.9 in increments of 0.1; therefore, each subject went through a sequence of 10 lottery choices. After all 10 choices were made, one lottery was randomly chosen for payment, although earnings from this task were revealed to subjects only at the very end of the experimental session.

[15] In the treatments where officials are allowed to post (BB_ALL and BB_IPB) we told subjects that both citizens and officials had access to the board and could post on it. In BB_CIT, we told subjects that officials could access the information on the board but only citizens could post on it.

[16] Natural exceptions to this were that we did not allow subjects to post messages containing office labels that were not on the map in the current round or bribes outside the [0, 50] range. We also did not allow a citizen to post more than one message about an office in each round. Thus, in each round a citizen can potentially post up to 7 messages about the 7 offices, regardless of how many offices he or she actually visited. The same restriction applies to the number of messages for officials (when they are allowed to post).

[17] While this aspect of a reporting system might be interesting to investigate, it does not really apply to the *I paid a bribe* website, hence we decided to abstract from it.

After all subjects completed the risk-aversion assessment task, the corruption experiment began. Subjects engaged in the experiment for 10 rounds,[18] and at the end of the session were paid the earnings based on their payoffs from one randomly selected round. Earnings from this part were converted from ECUs to US$ at the exchange rate of $1 for 15 ECU. The session concluded with a short questionnaire.

We conducted all experimental sessions at the XS/FS laboratory at Florida State University. The experiment was programmed in z-Tree (Fischbacher, 2007) and subjects were recruited among pre-registered FSU students using ORSEE (Greiner, 2015). In order to guarantee anonymity, at the beginning of each session subjects were randomly assigned an identification number, which they kept for the duration of the experiment. At no point did we ask subjects to reveal their names during the experiment, and although actual names were used during the payment process for accounting purposes, we informed the subjects that we would not register their names and, therefore, we would not be able to link them to the choices made in the experiment. Each session lasted between 60 and 90 minutes, with average earnings of $20.84 per subject including a $10 show-up fee.

3. Theoretical framework and predictions

Consider first the baseline treatment without information sharing. Assume that each public official $i = 1, \ldots, n$ has a privately observable prior, $\rho_i(b)$, about the distribution of bribes demanded by other officials. Similarly, each citizen $j = 1, \ldots, m$ has a privately observable prior, $\pi_j(b)$, about the bribes demanded by officials. The distributions of priors are common knowledge. At the beginning, the officials independently choose bribe levels $b_i \in [0, B]$. Then each citizen is randomly matched to an office and observes the bribe demanded by the official in that office. The citizen can either pay the bribe and get the license at that office or search, i.e., visit a different office and observe the bribe demanded there, at a search cost $c > 0$. The citizen can visit any number of offices, in addition to the office she is initially matched to, by paying the cost c for every new visit, but she has to obtain the license eventually, i.e., to pay the bribe at one of the visited offices. No search cost is incurred if the citizen returns to a previously visited office.[19]

The environment described above can be characterized formally as that of costly price search with recall. The most basic equilibrium prediction, going back to Diamond (1971), is that citizens will not search and all officials will charge the highest possible bribe B. This prediction, as well as other sequential and nonsequential search theories giving rise to price dispersion (see, e.g., Stahl, 1989; Burdett and Judd, 1983), rely on the assumption that citizens are aware of the distribution of bribe demands. We adopt a more realistic framework in which citizens gradually discover the distribution of bribes as they search. For such environments, it has been shown that the optimal search rule is sequential and myopic, i.e., in a given period the citizen decides whether or not to search further by comparing the expected gain from one search to the search cost c (see, e.g., Bikhchandani and Sharma, 1996).

Citizens' behavior. Suppose, without loss of generality, that citizen j is initially matched with office 1 and observes a bribe b_1. Her decision on whether or not to search depends on b_1 and on her beliefs $\pi_j(b|b_1)$ about the bribes in other offices, which is her updated prior. The citizen's expected gain from searching is $b_1 - \mathrm{E}_j(b|b_1)$, where $\mathrm{E}_j(\cdot\,|b_1)$ denotes the expectation, over $\pi_j(b|b_1)$, of a bribe demanded in the remaining offices. The citizen will search if that gain exceeds c, and will not search otherwise. Similarly, if the citizen searched through k offices (including the office she was initially matched to; assume, without loss of generality, that those are offices $1, 2, \ldots, k$), her gain from another search will be $\min\{b_1, \ldots, b_k\} - \mathrm{E}_j(b|b_1, \ldots, b_k)$, where $\mathrm{E}_j(\cdot|b_1, \ldots, b_k)$ denotes the expectation of a bribe demanded in the remaining $n - k$ offices with the beliefs updated by the observed bribes $b_1, \ldots, b_k$. Again, the citizen will search if that gain exceeds c, and will not search otherwise and acquire the license from the already visited office with the lowest demanded bribe.

Officials' behavior. We will assume that officials have moral costs associated with bribery. These moral costs can be due to aversion to corruption in general, aversion to violation of social norms, or due to fear of being observed by the experimenters or by other subjects. They may also be due to aversion to the inequality generated by bribery. Let $g_i(b;\ \rho_i(\ \cdot\ ))$ denote official i's moral cost of demanding bribe b. We assume that g_i is strictly increasing and convex in b, with $g_i(0;\rho_i(\cdot)) = 0$, and decreases in $\rho_i(\ \cdot\ )$ (in the FOSD sense) due to the effect of social norms. The official's expected utility from demanding bribe b then is $u_i(b) = K(b;\rho_i(\cdot))b - g_i(b;\rho_i(\cdot))$, where $K(b;\ \rho_i(\ \cdot\ ))$ is the expected number of citizens who will acquire the license at office i. The function $K(b;\ \rho_i(\ \cdot\ ))$ is decreasing in b and increasing in $\rho_i(\ \cdot\ )$ (in the FOSD sense).

Equilibrium in the bribery game can be defined as a configuration of bribes $b_1^*, \ldots, b_n^*$ such that every citizen follows the search rules described above and every official maximizes her utility $u_i(b_i)$.

Consider now how various information sharing systems can affect the equilibrium configuration of bribes. In general, the mere presence of a reporting system might act as a signal reinforcing the social norm that corruption is immoral and should be reported. This might shift downward citizens' and officials' priors about the distribution of bribes and hence shift upward officials' moral costs as compared to the baseline. However, besides this common feature of all corruption reporting

[18] Subjects were not informed about the total number of rounds in advance. Because of software glitches, two of the 24 sessions ended earlier: One BB_CIT session ended after 5 rounds, and one BB_IPB session ended after 7 rounds. Dropping these sessions does not qualitatively change the results.

[19] In the experiment, citizens also pay a fixed legal fee, F, and receive a fixed benefit, V, when they get the license, but we ignore these here because they have no effect on incentives.

mechanisms, different systems might have different effects on bribe demands depending on the kind of information they allow subjects to post and the restrictions they do or do not impose on the identity of the subjects who are allowed to post.

We start the analysis with the BB_CIT system, in which only citizens are allowed to post on the bulletin board, and the information includes both the size of the bribe and the office at which the bribe was demanded. Suppose, for simplicity, that every citizen posts truthfully the bribe she was asked for at the initial office. For citizens, such information about bribes completely eliminates the necessity to search beyond two times. Moreover, those who do search two times will guarantee themselves the lowest bribe.[20]

For officials, the information about bribes can have two effects. First, it allows them to form correct and symmetric beliefs about the bribes demanded by other officials. This information may either increase or decrease the moral costs, as compared to the case without information sharing, depending on how it changes social norms. Specifically, the officials whose prior beliefs about bribery were higher (lower) than the observed distribution of bribes may experience an increase (reduction) in moral costs. Second, the information about bribes exposes each official's bribe directly to all participants in the experiment. This exposure may have a direct shaming effect on the official and increase her moral costs. Assuming the initial beliefs about bribes demanded by other officials are, on average, correct, the overall effect of the BB_CIT information system on officials will be an increase in moral costs. In combination with the effect for citizens, this will lead to a reduction in equilibrium bribes.

A more realistic assumption on citizens' posting behavior – that only a certain fraction of citizens post on the bulletin board – will likely lead to the same comparative statics, although the effect will be weaker. It is possible that citizens will lie in their posts, thereby distorting the information about the distribution of bribes, although they have no obvious reasons to do so. It is also possible that the distribution of posted bribes will be different from the true distribution because the probability that a bribe is posted may depend on the size of the bribe, leading to nonrandom selection. For example, citizens may be more likely to post higher bribes because they want to punish the most corrupt officials. Alternatively, citizens may be more likely to post lower bribes because they want to help their fellow citizens and reward more honest officials. It is an empirical question which of the two effects will dominate, if any. However, due to the direct informational impact of the bulletin board we still expect bribery to be lower in the BB_CIT treatment as compared to the baseline.

Prediction 1. *Demanded and paid bribes will be lower in BB_CIT than in the baseline treatment.*

The information sharing system BB_ALL is different from BB_CIT in that public officials are also allowed to post on the bulletin board, and their posts are not distinguishable from those of the citizens. Unlike citizens, officials may have strong incentives to post false information.

(i) False advertisement about self: An official can post a low bribe for her own office hoping that citizens will visit it but then pay a higher bribe anyway because they decide not to search any more.

(ii) False advertisement about others: An official can post high bribes demanded in other offices to discourage citizens from visiting those offices. The goal is to attract visitors to her own office and/or to discourage search in general.

(iii) Information obfuscation and destruction of trust in the bulletin board: Official may decide to flood the bulletin board with false and contradictory information thereby obfuscating the truthful information posted by citizens and destroying the citizens' trust in the information sharing system altogether.

The effect of such an information system on bribery is ambiguous, and it is ultimately an empirical question. The strategic deterioration of information requires officials to lie, and a significant proportion of the population are averse to lying (e.g., Gneezy et al., 2013). Overall, because the quality of information will likely be reduced strategically at least by some officials, we expect higher average bribes in BB_ALL as compared to BB_CIT. It is not clear, however, if BB_ALL will lead to lower bribes than the baseline.

Prediction 2. *Demanded and paid bribes will be higher in BB_ALL than in BB_CIT.*

The information sharing system BB_IPB is intended to mimic the *I paid a bribe* information system. It is different from BB_ALL in that citizens and officials can only post the size of the bribes demanded in various offices but cannot post information about the office where the bribe was demanded. Thus, even though the (possibly distorted) information about the distribution of bribes is still available, citizens can no longer use the information system to narrow down their search, and direct shaming of individual officials is impossible. At the same time, officials can no longer engage in specific false advertisements about themselves and others, and will only post false information in order to distort the true distribution of bribes. Therefore, we expect a weaker deterioration of information as compared to BB_ALL, but the resulting (more truthful) information will be less powerful than in BB_CIT in terms of helping citizens locate lower bribes. The effect on officials' moral costs is also ambiguous, as compared to BB_ALL, because more realistic beliefs about the distribution of bribes may increase moral costs, but the absence of direct shaming may reduce them. Overall, while we still predict higher average bribes in BB_IPB than in BB_CIT, it is not clear how BB_IPB will compare to BB_ALL or to the baseline.

Prediction 3. *Demanded and paid bribes will be higher in BB_IPB than in BB_CIT.*

[20] Recall that in the experiment citizens are only allowed to observe the bulletin board after they have made their first search decision. The second search decision in this case does not help discover new information but simply allows the citizen to visit the office with the lowest bribe if it has not already been visited. This decision still requires paying the search cost.

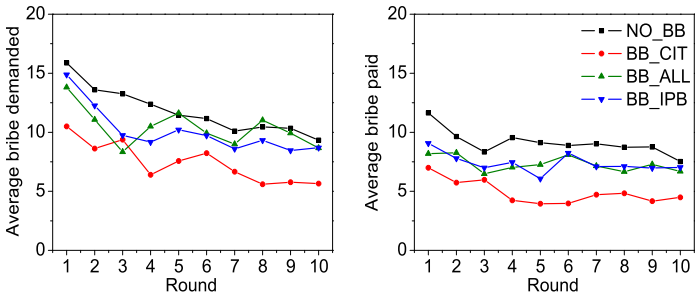

Fig. 1. Average bribe demanded by officials (left) and paid by citizens (right) in each round, by treatment.

Table 2
Summary statistics by treatment.[a]

	Bribe demanded		Bribe paid		Citizens' earnings	
	All t	$t > 5$	All t	$t > 5$	All t	$t > 5$
NO_BB	11.80	10.28	9.12	8.59	118.85	119.72
	(0.60)	(0.52)	(0.64)	(0.70)	(0.70)	(0.63)
BB_CIT	7.54	6.38	4.95	4.43	123.34	124.45
	(0.60)	(0.41)	(0.35)	(0.33)	(0.32)	(0.14)
BB_ALL	10.39	9.71	7.31	7.18	121.17	121.49
	(0.81)	(0.54)	(0.76)	(0.47)	(0.85)	(0.52)
BB_IPB	10.17	8.97	7.40	7.32	120.42	120.72
	(1.06)	(0.62)	(0.73)	(0.72)	(0.95)	(0.84)

[a] Standard errors in parentheses are clustered by session.

4. Results

4.1. Treatment effects

Fig. 1 shows average bribes demanded by officials (left) and paid by citizens (right) in each round, by treatment. As seen from the figure, the rankings of bribes demanded and paid across treatments are similar, although, as expected, bribes demanded are higher than the bribes ultimately paid. Bribes demanded have a more pronounced downward trend, especially at the beginning, whereas bribes paid appear relatively more stable over time. Fig. 1 suggests that average bribes demanded are the lowest in the BB_CIT treatment, while the other three treatments are difficult to rank. Bribes paid appear to be the highest in the baseline treatment (NO_BB) and the lowest in BB_CIT, with the other two treatments (BB_ALL and BB_IPB) in between.

Table 2 reports average bribes demanded and paid using data from all 10 rounds (labeled "all t") and from the second half of the experiment (rounds 6–10, labeled "$t > 5$"). As seen from the table, average bribes demanded and paid in BB_CIT are lower than in all other treatments. Compared to the baseline, bribes demanded in BB_CIT are nearly 40% lower, and bribes paid are nearly 50% lower. Bribes in BB_ALL and BB_IPB are below the baseline too, but in many cases the differences are small, and there are no systematic differences between BB_ALL and BB_IPB. The observed differences in citizens' earnings, also reported in Table 2, reflect the differences in bribes, with the highest earnings in BB_CIT and lowest in the baseline treatment.

For the most basic comparison of bribes demanded and paid across treatments, we performed pairwise Wilcoxon rank-sum tests comparing average bribes demanded and paid in all rounds, and in rounds $t > 5$, treating each session as one independent observation. In all cases, bribes demanded and paid in BB_CIT are significantly lower as compared to each of the other three treatments, but all other comparisons do not produce statistically significant differences.[21]

Table 3 presents the results of basic OLS regressions, with standard errors clustered at the session level, testing for differences in average bribes demanded and paid across treatments. The regressions include treatment dummies *BB_CIT*, *BB_ALL* and *BB_IPB* (with the NO_BB treatment as the baseline), and the time trend (*Round*). Similar to Table 2, we present the results using data from all 10 rounds ("all t") and rounds 6–10 ("$t > 5$").[22] The regression results show, consistent with the nonparametric tests, that average bribes demanded and paid are significantly lower in BB_CIT as compared to the baseline in all instances. The Wald tests comparing the coefficients on *BB_CIT*, *BB_ALL* and *BB_IPB* confirm that bribes

[21] Average bribe demanded is lower in BB_CIT as compared to NO_BB ($p = 0.0065$ for all t, $p = 0.0062$ for $t > 5$), to BB_ALL ($p = 0.025$ for all t, $p = 0.0061$ for $t > 5$) and to BB_IPB ($p = 0.037$ for all t, $p = 0.0062$ for $t > 5$). Average bribe paid is lower in BB_CIT as compared to NO_BB ($p = 0.0039$ for all t, $p = 0.0062$ for $t > 5$), to BB_ALL ($p = 0.025$ for all t, $p = 0.0062$ for $t > 5$) and to BB_IPB ($p = 0.037$ for all t, $p = 0.0029$ for $t > 5$).

[22] Recall that due to a software problem the experiment ended after 5 rounds in one of the sessions; therefore, the number of clusters (sessions) is 23 for $t > 5$.

Table 3
OLS regression results for average treatment effects.[a]

	Bribe demanded		Bribe paid	
	All t	$t > 5$	All t	$t > 5$
BB_CIT	−4.37***	−3.89***	−4.21***	−4.15***
	(0.82)	(0.62)	(0.69)	(0.72)
BB_ALL	−1.41	−0.57	−1.81*	−1.41*
	(0.94)	(0.70)	(0.93)	(0.79)
BB_IPB	−1.71	−1.34*	−1.75*	−1.29
	(1.06)	(0.74)	(0.89)	(0.94)
Round	−0.48***	−0.32**	−0.19***	−0.20**
	(0.09)	(0.14)	(0.07)	(0.08)
Constant	14.46***	12.85***	10.15***	10.21***
	(0.85)	(1.35)	(1.19)	(0.96)
Observations	1624	784	1624	784
Subjects	168	161	168	161
Clusters	24	23	24	23
R^2	0.068	0.058	0.093	0.122

[a] Standard errors clustered by session are shown in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

demanded and paid in BB_CIT are also lower than in BB_ALL ($p = 0.0057$ for all t, $p = 0.000$ for $t > 5$ for bribes demanded; $p = 0064$ for all t, $p = 0.000$ for $t > 5$ for bribes paid) and in BB_IPB ($p = 0.023$ for all t, $p = 0.001$ for $t > 5$ for bribes demanded; $p = 0033$ for all t, $p = 0.001$ for $t > 5$ for bribes paid), and there are no statistically significant differences between BB_ALL and BB_IPB. Bribes demanded in BB_ALL are not different from those in the baseline, while bribes paid are somewhat lower. Importantly, in the *I paid a bribe* treatment BB_IPB bribes demanded and paid are not systematically lower than those in the baseline. We summarize the results as follows.

Result 1. *The reporting system that restricts posting to citizens only and requires specific information about reported offices (BB_CIT) is the only system that systematically and significantly reduces bribe demands as compared to the no-reporting baseline, and is the most effective in reducing bribe payments.*

Result 2. *The reporting system that allows both citizens and public officials to post and requires specific information about reported offices (BB_ALL) does not reduce bribes demanded, but it does reduce bribes paid as compared to the no-reporting baseline.*

Result 3. *The reporting system that mimics the "I paid a bribe" website, in which both citizens and public officials can post without providing specific information about bribe-demanding offices (BB_IPB), does not systematically reduce bribes demanded and paid as compared to the no-reporting baseline.*

Fig. 2 shows the histograms of bribes demanded by treatment. The bins on the horizontal axes are bribe intervals [0, 4], [5, 9], [10, 14], ..., [45, 50], and the tick labels indicate the left end of each bin.[23] As seen from the figure, all three BB treatments produce higher frequencies of bribes in the lowest two bins (30.7%, 18.1%, 16.3% in the first bin and 39.2%, 27.6%, 29.6% in the second bin for BB_CIT, BB_ALL and BB_IPB, respectively) as compared to the baseline (10% in the first bin and 17.4% in the second bin). In treatment BB_CIT the shift is the most pronounced, with almost 70% of bribes demanded falling in the first two bins (9 or lower). In BB_ALL and BB_IPB the distributions are more spread out, with higher frequencies of bribes in the third bin and above as compared to BB_CIT.

4.2. Dynamics

Table 4 shows the results of dynamic regressions for bribes demanded by officials and bribes paid by citizens. In addition to the treatment dummies and time trend, we control for the bribe demanded (for officials) and paid (for citizens) in the previous round ($Bribe_{t-1}$ and $Bribe\_paid_{t-1}$, respectively). Additionally, for officials we control for the number of citizens who decided to get the license from the official in the previous round, normalized by the total number of the office visitors ($\%Bought_{t-1}$), and for the interaction of this variable with the bribe demanded in the previous round. For citizens, we control also for the number of times the citizen searched in the current round (*# ofSearches*). As a robustness check, we report the results for all periods (All t) and periods 6–10 ($t > 5$) including and not including demographic factors (the subject's gender, age and a dummy equal to 1 if he or she is an economics major).

As expected, due to persistence in bribe demands, the coefficients on the treatment dummy variables become smaller in magnitude as compared to those reported in Table 3, and only the ones on *BB_CIT* remain statistically significant. Persistence in bribe demands increases in the proportion of citizens who decided to obtain the license from the official in the previous

[23] About 75% of bribes demanded are multiples of 5. We chose the first bin not to include 5 because payoff-maximizing citizens encountering a bribe lower than 5 should not search further. This indeed occurs in 90.6% of the cases.

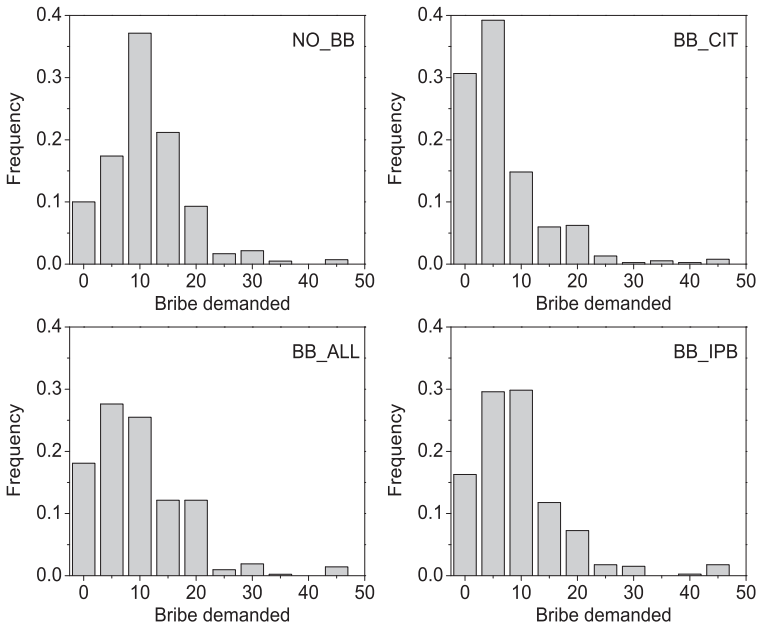

Fig. 2. The distribution of bribe demands, by treatment. The bins on the horizontal axes are bribe intervals [0, 4], [5, 9], [10, 14], . . ., [45, 50].

Table 4
OLS regression results for individual-level decisions on bribes demanded.[a]

	Bribe demanded (officials)			Bribe paid (citizens)		
	All t	$t > 5$	$t > 5$	All t	$t > 5$	$t > 5$
BB_CIT	−1.78***	−1.47***	−1.55***	−3.24***	−3.40***	−3.57***
	(0.54)	(0.38)	(0.36)	(0.58)	(0.66)	(0.62)
BB_ALL	−0.49	−0.14	−0.24	−1.26*	−1.16*	−1.24*
	(0.41)	(0.31)	(0.29)	(0.72)	(0.62)	(0.60)
BB_IPB	−0.68	−0.33	−0.29	−1.20	−0.90	−0.88
	(0.48)	(0.41)	(0.42)	(0.71)	(0.77)	(0.77)
$Bribe_{t-1}$	0.43***	0.42***	0.43***			
	(0.07)	(0.09)	(0.09)			
$Bribe_{t-1} \times Bought_{t-1}$	0.15***	0.17***	0.17***			
	(0.03)	(0.05)	(0.05)			
$\%Bought_{t-1}$	−0.38	−0.98	−1.01			
	(0.60)	(1.24)	(1.22)			
Round	−0.077	−0.12	−0.12	−0.073	−0.21**	−0.21**
	(0.060)	(0.11)	(0.11)	(0.045)	(0.08)	(0.08)
$Bribe\_paid_{t-1}$				0.22***	0.19***	0.18***
				(0.03)	(0.06)	(0.06)
# of Searches				−0.58***	−0.71***	−0.69***
				(0.14)	(0.13)	(0.12)
Constant	5.47***	5.89***	5.25***	7.42***	8.85***	8.56***
	(1.29)	(1.58)	(1.75)	(0.68)	(0.98)	(1.14)
Demographics	No	No	Yes	No	No	Yes
Observations	1456	784	784	1456	784	784
Subjects	168	161	161	168	161	161
Clusters	24	23	23	24	23	23
R^2	0.311	0.350	0.353	0.164	0.172	0.180

[a] Standard errors clustered by session are shown in parentheses. Significance levels: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

round (as evidenced by the positive and significant coefficients on the interaction $Bribe_{t-1} \times \%Bought_{t-1}$). From the results of the regression for $t > 5$ including demographics, a 10 percentage points increase in $\%Bought_{t-1}$ leads to a 0.017 increase in the persistence. The effect of a 10 percentage points increase in $\%Bought_{t-1}$ on the level of bribes in the current period is $-0.101 + 0.017Bribe_{t-1}$, i.e., it is positive for bribes above 6.

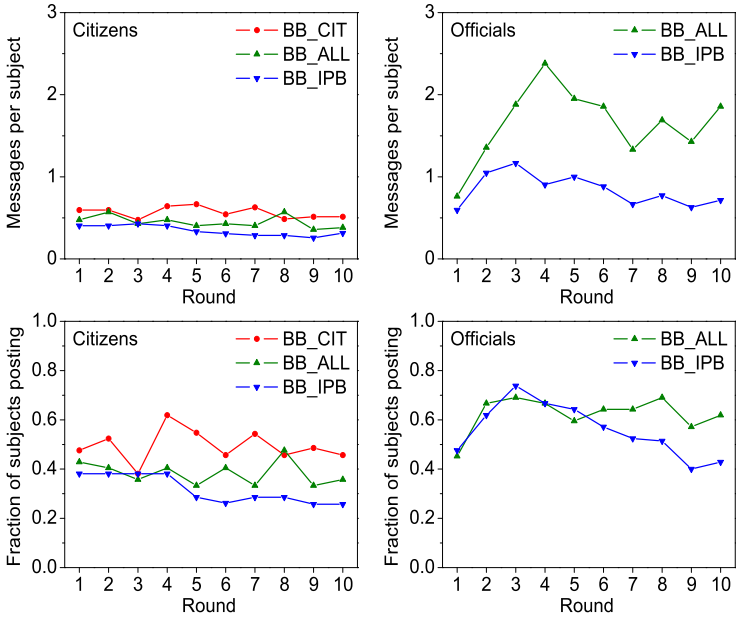

Fig. 3. *Upper panels*: The average number of messages per subject posted by citizens (left) and by officials (right) in a given round, by treatment. *Lower panels*: The fraction of citizens (left) and officials (right) who posted on the bulletin board at least once in a given round, by treatment.

For citizens, the persistence of bribes paid is not as strong, especially in the later rounds, and hence the treatment effects of BB_CIT and BB_ALL are still statistically significant in the dynamic models, similar to Table 3. The number of searches affects bribes paid negatively, as expected.

4.3. Messages and lies

We now turn to the analysis of messages posted by citizens and officials in the three BB treatments. Recall that citizens were allowed to post messages in all three treatments, while officials could only post in BB_ALL and BB_IPB. The top two panels in Fig. 3 show the average number of messages posted, per subject, in each round by citizens (upper left) and officials (upper right). On average, there are 0.57, 0.45 and 0.35 messages per subject per period posted by citizens in BB_CIT, BB_ALL and BB_IPB, respectively. Pairwise comparisons of the average number of citizens' posts between treatments via a regression on treatment dummies, with standard errors clustered by session, produce a statistically significant difference between BB_CIT and BB_IPB ($p = 0.000$; standard errors clustered by session), but not between BB_CIT and BB_ALL ($p = 0.266$) or between BB_ALL and BB_IPB ($p = 0.332$). Officials post substantially more than citizens, when given a chance – 1.65 ($p = 0.003$) and 0.84 ($p = 0.000$) messages per subject per period in BB_ALL and BB_IPB, respectively. Moreover, officials post more messages in BB_ALL and BB_IPB even compared to citizens in BB_CIT ($p = 0.006$ and $p = 0.007$, respectively). The difference between BB_ALL and BB_IPB for officials is statistically significant ($p = 0.058$). Thus, both citizens and officials post less in BB_IPB as compared to other BB treatments. This is likely due to the fact that bribery information is not specific in BB_IPB and hence posting is less effective.

The lower two panels in Fig. 3 shows the fraction of citizens (left) and officials (right) who posted at least once in a given round, by treatment. On average, 49.6%, 38.3% and 31.8% of citizens posted at least once per round in BB_CIT, BB_ALL and BB_IPB, respectively. Similar to the overall posting rates, the difference between BB_CIT and BB_IPB is statistically significant ($p = 0.007$), while the differences between BB_CIT and BB_ALL, and between BB_ALL and BB_IPB are not ($p = 0.182$ and $p = 0.454$, respectively). For officials, the averages are higher – 62.4% in BB_ALL ($p = 0.062$) and 56.4% in BB_IPB ($p = 0.000$). The difference between BB_ALL and BB_IPB is not significant ($p = 0.564$). Thus, citizens are more likely to leave messages on the board when they can post specific information and when they know that officials cannot post on the board. Officials, when given a chance to post (in BB_ALL and BB_IPB), are more likely than citizens to post at least once.[24]

[24] As rightly noted by a reviewer, one reason why officials may be posting more than citizens in the experiment is that they have nothing else to do while waiting for citizens to make their search and paying decisions. However, if boredom were the main driver of officials' posting behavior, we would see the same or even higher number of messages posted by officials in BB_IPB as compared to BB_ALL (recall that it is easier to post a message in BB_IPB because one only needs to type in the bribe amount, whereas in BB_ALL one also needs to type in the office number). Instead, we observe substantially more postings by officials in BB_ALL (cf. Fig. 3, upper right). This suggests that there is a significant strategic component in the motives behind the decisions by officials to post on the board, and, although boredom may be a reason for posting, it is not the only reason.

Table 5
Summary statistics for messages and lies, by treatment, aggregated over all rounds.[a]

	Citizens			Officials		
Treatments	% who posted	Avg. # of posts	% false posts	% who posted	Avg. # of posts	% false posts
BB_CIT	49.6	1.15	5.8			
	(3.9)	(0.05)	(1.8)			
BB_ALL	38.3	1.17	4.5	62.4	2.65	54.4
	(7.7)	(0.05)	(1.7)	(10.4)	(0.28)	(6.1)
BB_IPB	31.8	1.09	12.2	56.4	1.50	92.7
	(4.8)	(0.04)	(2.7)	(4.9)	(0.05)	(1.5)

[a] Standard errors clustered by session are in parentheses. The average numbers of posts are restricted to the citizens or officials who posted at least once.

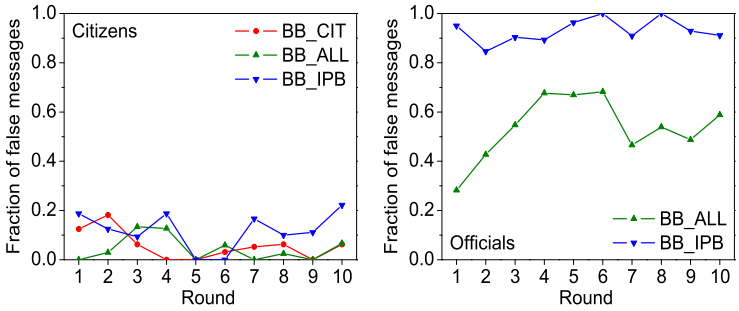

Fig. 4. The percentage of false messages posted by citizens (left) and officials (right), by treatment.

Table 5 provides additional summary statistics focusing only on those citizens and officials who posted on the bulletin board at least once in a given round. For these citizens and officials, Table 5 shows the average number of posts and the percentages of false messages in each treatment. Interestingly, although fewer citizens post messages in BB_IPB, those who do post produce about the same average number of messages as in BB_CIT and BB_ALL. Thus, the difference in citizens' posting rates between treatments is at the extensive margin only. This is not the case for officials who are equally likely to post in BB_IPB as compared to BB_ALL, but generate fewer messages (1.50 versus 2.65 per subject, $p = 0.002$) if they decide to post.

How truthful is the information posted by citizens and officials? Fig.4 shows the proportion of false messages posted by citizens (left) and officials (right). In BB_CIT and BB_ALL, a message is defined as false for a citizen if she posts a bribe/office combination she did not encounter so far in the current round. In BB_IPB, a citizen's post is defined as false if the citizen posts a bribe she has not encountered so far in the current round. For an official, in BB_ALL a message is defined as false if it refers to an office other than the official's own office (since officials have no information about bribes demanded by other officials), or if it is about the official's own office but reports a bribe that is different from what the official demanded. In BB_IPB, any post referring to a bribe different from the official's own is considered false.

As seen from Table 5 and Fig. 4, when posts contain specific information about the location of the offices demanding bribes (BB_CIT and BB_ALL) citizens are very unlikely to lie in their posts. Indeed, both in BB_CIT and BB_ALL less than 6% of the messages posted by citizens contain false information. Interestingly, the percentage of deceptive posts by citizens increases to 12.2% ($p = 0.049$ comparing BB_IPB to BB_CIT and $p = 0.020$ comparing BB_IPB to BB_ALL) in the *I paid a bribe* treatment.[25] In contrast, lying by officials is widespread, especially in *I paid a bribe* where 92.7% of officials' posts contain false information (versus 54.4% in BB_ALL, $p = 0.000$). Thus, as predicted, when given a chance, officials try to strategically distort the information on the bulletin board. In what follows, we explore in more detail the types of false information posted by officials in treatment BB_ALL where in addition to the size of the bribe officials could post the location where the bribe was supposedly encountered.

The left panel in Fig. 5 shows the fraction of officials in treatment BB_ALL who lied about self, about other officials, or both, in each round. As seen from the figure, most lying officials lie about both self and others, utilizing the whole set of available obfuscation strategies. The right panel in Fig. 5 shows the fraction of officials lying in BB_ALL for different levels of bribes they demanded (N is the number of observations in each bribe range). As seen from the figure, officials who demand

[25] Overall, the average bribe posted by citizens is slightly (1.77 ECU), but significantly ($p = 0.000$) above the actual average bribe, and there are no significant differences in this number across treatments. This may be partly due to the nonrandom selection of posted bribes, i.e., citizens being more likely to post higher bribes, and partly due to lying. Restricting the sample to those citizens who posted truthful messages, we obtain the difference between the average posted bribe and the actual average bribe of 1.22. On the other hand, restricting the sample to those citizens who lied in their messages produces the differences of 12.47, 10.37 and 5.32 in BB_CIT, BB_ALL and BB_IPB, respectively. Thus, the selection effect on the bribes posted by citizens is relatively minor, but when citizens lie they tend to post much higher bribes than those they observe.

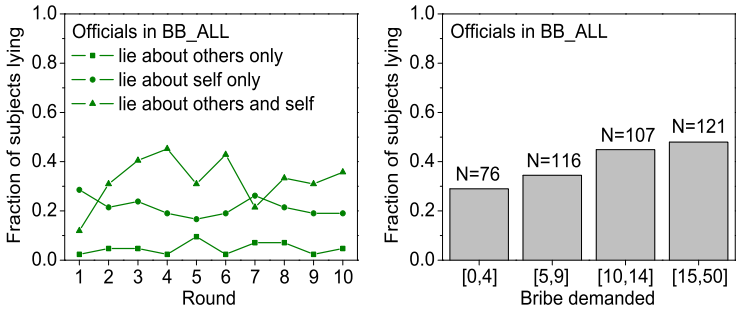

Fig. 5. *Left*: The fraction of officials who lied about self, others or both in BB_ALL, by round. *Right*: The fraction of officials who posted false messages by the size of the bribe they demanded in BB_ALL. *N* is the number of observations in each bribe range.

higher bribes also tend to lie more in their messages. Only 28.9% of officials who demand bribes in the [0, 4] range lie, as compared to 47.9% of officials who demand bribes of 15 and higher. A probit regression of the individual probability of lying on bribe demanded, restricted to officials who post a positive number of messages, produces a marginal effect of 0.017 ($p = 0.000$), i.e., a 5 ECU increase in the bribe demanded is associated with an 8.5 percentage point increase in the probability of lying.

The results can be summarized as follows.

Result 4. *(a) When specific information is required in posts (in treatments BB_CIT and BB_ALL), between 40% and 50% of the citizens post on the board, and lying by citizens is unlikely.*

(b) When specific information is not required (in treatment BB_IPB), fewer citizens post on the board and lying by citizens is more frequent, albeit still relatively unlikely.

Result 5. *(a) When allowed, posting by officials is widespread. Officials post more when specific information is required in posts (in treatment BB_ALL) as compared to when it is not required (in treatment BB_IPB).*

(b) Most officials post false information, and most of the messages posted by officials are false. Officials lie more when specific information is not required in posts (in treatment BB_IPB) as compared to when it is required (in treatment BB_ALL).

5. Conclusions

Thanks to the efforts of non-profit organizations like Transparency International, as well as media outlets and academic publications, the awareness of corruption, its causes and consequences, has dramatically increased in the last two decades. Despite these important developments, examples of successful anti-corruption strategies are hard to find. Identifying policies that might effectively reduce corruption is challenging, especially in countries that are systemically corrupt, where acting on top-down monitoring and enforcement may easily backfire, due to the existence of well-ingrained systems of bribe and favor exchanges between different layers of government bureaucracies. In these societies, bottom-up mechanisms, such as crowdsourced reporting platforms, might represent the only viable alternative. While these platforms have appeared in many highly corrupt countries and they are rapidly spreading around the globe, their goal so far has been merely to raise awareness about the existence and frequency of corrupt exchanges. The role that these platforms could play in actively fighting corruption is still unknown due to both their nation-wide, rather than geographically randomized, implementation and the difficulty of quantifying corruption before and after their appearance.

In this paper, we contributed to the anti-corruption debate by experimentally testing the effectiveness of a corruption reporting platform resembling the *I paid a bribe* website first introduced in India in 2010. In an attempt to investigate whether the impact of crowdsourced reporting platforms on corruption could be enhanced, we conducted two additional treatments simulating modified versions of the existing reporting websites. The underlying criterion we were adamant to maintain when designing these additional platforms was the absence of any top-down monitoring and probabilistic punishment. Instead, we asked whether bribe demands would be lower if we could transform existing corruption reporting platforms into information sharing mechanisms that would allow citizens to search for and easily identify low bribe-demanding service providers.

Like the *I paid a bribe* website, we focused on extortionary corruption, i.e., the demand and payment of harassment bribes for the provision of goods or services that citizens are entitled to. While this kind of corruption is less likely to make news headlines than corruption scandals involving high-level officials and large amounts of money, it permeates every-day life in most developing countries, generating high financial burdens and often preventing the most vulnerable segments of the population from accessing basic services.

In our laboratory experiment, we allowed subjects in the role of public officials to demand harassment bribes for the provision of a license. Subjects in the role of citizens could pay the requested bribe or visit a different office. In the treatment simulating the *I paid a bribe* website (BB_IPB), we allowed citizens to post messages reporting the demanded bribes on a Bulletin Board (BB). Like in the *I paid a bribe* website, the messages were anonymous and did not report the exact location

of the office where the reported bribe was demanded. Moreover, access to the BB was not restricted to citizens, i.e., officials were also allowed to post on the board. This way, we were able to assess to what extent the effectiveness of corruption reporting may be reduced by the presence of false and non-representative information, including strategic lies by public officials. In our two other BB treatments, the messages posted on the board, albeit still anonymous, now also contained information about the exact location of the offices where the reported bribes were demanded. The difference between these two additional treatments lies only in the presence or absence of restrictions concerning the identity of the subjects allowed to post on the board. In one treatment (BB_CIT) we restricted posting to citizens, while in the other (BB_ALL) we also allowed officials to post. In all BB treatments we allowed for the posting of false information.

Our experimental results are insightful with respect to both the existing *I paid a bribe* website and a new generation of corruption reporting platforms. Our findings suggest that the possibility to post on the *I paid a bribe* website may be insufficient to significantly and systematically lower the occurrence and the size of harassment bribes. This is not very surprising, since the purpose of the *I paid a bribe* website was never to act as a policy instrument in the actual fight against corruption. Our second finding is that making the information on the reports more specific by allowing for the identification of the location of the bribe exchange, while maintaining anonymity, significantly reduces the bribes paid by citizens. This result holds even if posting on the platform is open to any member of the public, including public officials, and if, as observed in the experiment, the utilization rate by citizens is relatively low as compared to that of officials. Finally, the most effective platform is one where the information posted is specific about the location of the bribe-demanding official and where posting is restricted to citizens. This is because, contrary to officials, citizens rarely lie in their posts.

Overall, our investigation generates three important messages. First, the presence of a crowdsourced corruption reporting platform like the *I paid a bribe* website does not lead to systematically lower bribery as compared to a setting where no bottom-up reporting is allowed. Second, harassment bribes could be reduced if citizens' reports contained specific information about the location of the office/official where the bribes were demanded, transforming the platform into a crowdsourced search engine for low bribe-demanding officials. The website and app Bribespot is an example of how this could be achieved by using the location of the individuals reporting bribery to create corruption "hot spots" on a map easily accessible to other users.[26] Third, restricting the posting on the platform to service recipients would be valuable for the achievement of better anti-corruption outcomes. Without such restrictions, we found posting and lying by officials on the platform to be widespread and problematic for the effectiveness of the platform, especially in the long run. Since in our setting we had an equal number of citizens and officials, posts by officials are likely more widespread in our experiment than they could be in the field.[27] Nevertheless, our results highlight that preventing officials from posting on the reporting websites is a very desirable feature of any crowdsourced platform actively engaged in the fight against corruption.

A successful corruption reporting platform should preserve anonymity of posts but simultaneously should be able to uniquely identify and verify each poster's credentials. This could be achieved differently in different countries depending on the extent to which corruption does or does not permeate the highest levels of government. In settings where top government officials are truly committed to the fight against bureaucratic corruption, the reporting platform could rely on unique identifiers only visible to service recipients and highest level officials. In settings where government's anti-corruption efforts are poor and/or not credible, the identification process could be outsourced to a trustworthy third party, possibly an international organization.

Appendix A. Experimental instructions and screens (treatment BB_ALL)

General instructions

Thank you all for coming today. You are here to participate in an experiment. After playing the game you will be asked to complete a brief questionnaire. In addition to a $10 participation fee, you will be paid any money you accumulate from the experiment. You will be paid privately, by check, at the conclusion of the experiment. This study has been reviewed and approved by the FSU Human Subjects Committee. If you have any questions during the experiment, please raise your hand and wait for an experimenter to come to you. Please do not talk, exclaim, or try to communicate with other participants during the experiment. Participants intentionally violating these rules may be asked to leave the experiment and may not be paid.

Please read and sign the Consent form that you found on your desk. Please raise your hand if you have any question about any of the information on the Consent form. We will proceed with the experiment once we have collected all signed consent forms.

The number that you have found on your desk is your identification number in the experiment. We won't ask you to write down your name at any time during this experimental session. No one, including the experimenter, will have a way to link your name to the decisions you made in the experiment. At the end of the session, you will need to show your number to the experimenter in order to receive the money that you collected in the experiment.

[26] See: http://bribespot.com/en/. We thank a reviewer for bringing this website to our attention. One limitation of Bribespot.com is that it is still not possible to fully use the website as a search engine when in need of a specific service (say a driver's license). In other words, the corruption hot spots in the map are based on aggregated reports on all kinds of services for which different officials demanded bribes in a given location.

[27] A caveat here is that it is possible for public officials to recruit trolls or use automated posting programs (bots) in order to distort information on anti-corruption boards.

Earnings during the experiment will be denominated in Experimental Currency Units, or ECU. At the end of the experiment your earnings will be converted to dollars at the exchange rate of $1 for 15 ECU.

The experiment will consist of several parts and the instructions will be provided separately at the beginning of each part.

Instructions for Part 1

In each round of this series of decisions you will be asked to make a choice between two lotteries that will be labeled A and B. There will be a total of 10 rounds and after you have made your choice for all 10 rounds, one of those rounds will be randomly chosen to be played. Lottery A will always give you the chance of winning a prize of $2.00 or $1.60, while lottery B will give you the chance of winning $3.85 or $0.10. Each decision round will involve changing the probabilities of your winning the prizes. For example in round 1, your decision will be represented on the screen in front of you:

Your decision is between these two lotteries:

Lottery A: A random number will be drawn between 1 and 100. You will win

$1.60 if the number is between 1 and 90 (90% chance)

$2.00 if the number is between 91 and 100 (10% chance)

Lottery B: A random number will be drawn between 1 and 100. You will win

$0.10 if the number is between 1 and 90 (90% chance)

$3.85 if the number is between 91 and 100 (10% chance)

If you were to choose lottery B and this turns out to be the round actually played, then the computer will generate a random integer between 1 and 100 with all numbers being equally likely. If the number drawn is between 1 and 90, then you would win $0.10 while if the number is between 91 and 100, then you would win $3.85. Had you chosen lottery A, then if the number drawn were between 1 and 90 you would win $1.60 while a number between 91 and 100 would earn you $2.00.

All of the other 9 choices will be represented in a similar manner. Each will give you the probability of winning each prize as well as translate that probability into the numerical range the random number has to be in for you to win that prize.

At the end of the 10 choice rounds, the computer will randomly pick one of the 10 rounds to base your payment on, and draw the random number between 1 and 100 to determine your earnings. You will be informed about your earnings from this part of the experiment at the very end after you complete all parts.

Are there any questions before you begin making your decisions?

You will now start the sequence of 10 choices. You will be able to go through the choices at your own pace, but we will not be able to continue the experiment until everyone has completed this series.

Instructions for Part 2

This part of the experiment will consist of several decision sequences. The instructions will be given separately at the beginning of each sequence. At the end of the experiment one of the sequences will be randomly chosen to base your actual earnings on.

SEQUENCE 1

You are going to participate in this experimental task in one of two possible roles. You will be randomly assigned either the role of Public Official or the role of Private Citizen. A total of 7 Public Officials and 7 Private Citizens will participate in the task.

Each Public Official will be in charge of an Office that provides licenses to Private Citizens, and will receive a lump-sum wage of 130 ECU.

Each Private Citizen will start with a monetary endowment of 80 ECU and will have to get a license from one of the 7 Offices. The license will generate a monetary benefit of 70 ECU to the Citizen. The Private Citizen will have to pay a fee in order to get the license. The official license fee is 20 ECU. However, Public Official can refuse to provide the license unless a bribe is paid on top of the official fee. The bribe demanded by a Public Official can be any integer amount between 1 and 50 ECU.

At the beginning, each Public Official will decide whether or not to demand a bribe from the Private Citizens who may visit his or her Office, and the specific amount of the bribe, in the range between 1 and 50 ECU. The decision to demand a bribe and the size of the bribe cannot be changed during the sequence.

Each Public Official will not know if the other Public Officials chose to demand a bribe or the size of their bribes, if any. Private Citizens will also be initially unaware of the bribes demanded by each Public Official, if any, but they will be able to acquire such information by visiting the corresponding Office, at the cost of 5 ECU for every new visit.

The sequence proceeds as follows:

- At the beginning, each Public Official has to decide whether he or she would like to request a bribe, between 1 and 50 ECU, for the provision of the license, on top of the official fee of 20 ECU.

- Each Private Citizen is initially randomly assigned to visit an Office and finds out if a bribe is requested by the Public Official in that Office, and if so, the size of the bribe.

- Then, the Private Citizen has to decide whether to pay the total amount requested by the visited Office, and receive the license there, or leave that Office and choose to visit any of the other 6 available Offices. Every visit to a new Office costs 5 ECU to the Private Citizen.

- The Citizen can visit as many Offices as he or she wishes, at the cost of 5 ECU for any new visit, and can acquire the license from any of the Offices previously visited by paying the amount requested by the Official in that Office.

- The Private Citizen has to get the license eventually.

The payoffs from the sequence are determined as follows:

• Each Public Official earns a lump-sum wage of 130 ECU. On top of the wage, if the Public Official decides to demand a bribe for his or her services, he or she can get additional earnings from the bribes paid by the Private Citizens who visited the Office and decided to obtain the license there, if any.

• Each Private Citizen starts with an endowment of 80 ECU. When the Private Citizen gets the license, he or she additionally receives 70 ECU, but will have to pay the total amount requested by the Public Official (that may or may not include a bribe) and the accumulated cost of office visits, which is equal to 5 ECU × (number of visited offices).

Private Citizens will see the map below, showing the available 7 Offices that they can visit to get the license. By clicking on an Office, Private Citizens will be able to visit that Office and get information about whether a bribe is requested by the corresponding Public Official, and the size of the bribe, if any.

After being initially matched with one Office, each Citizen will be able to visit as many Offices as he or she wishes, at the cost of 5 ECU per new visit. Once an Office has been visited, Citizens will be able to see the requested amount on the map, in the corresponding box. Citizens could decide to get the license from any of the Offices previously visited, or visit a new Office.

After visiting an Office, Citizens will be able to access a Bulletin Board where they will have the chance to post messages about the bribes demanded by the officials in the visited Offices, if any. Citizens can access and post on the Bulletin Board at no cost. Officials will also be able to access the board and post messages at no cost to them. By accessing the Board, all subjects will see the messages previously posted by themselves and by others, if any. Whether participants will leave messages on the Board or not is up to them.

Are there any questions?

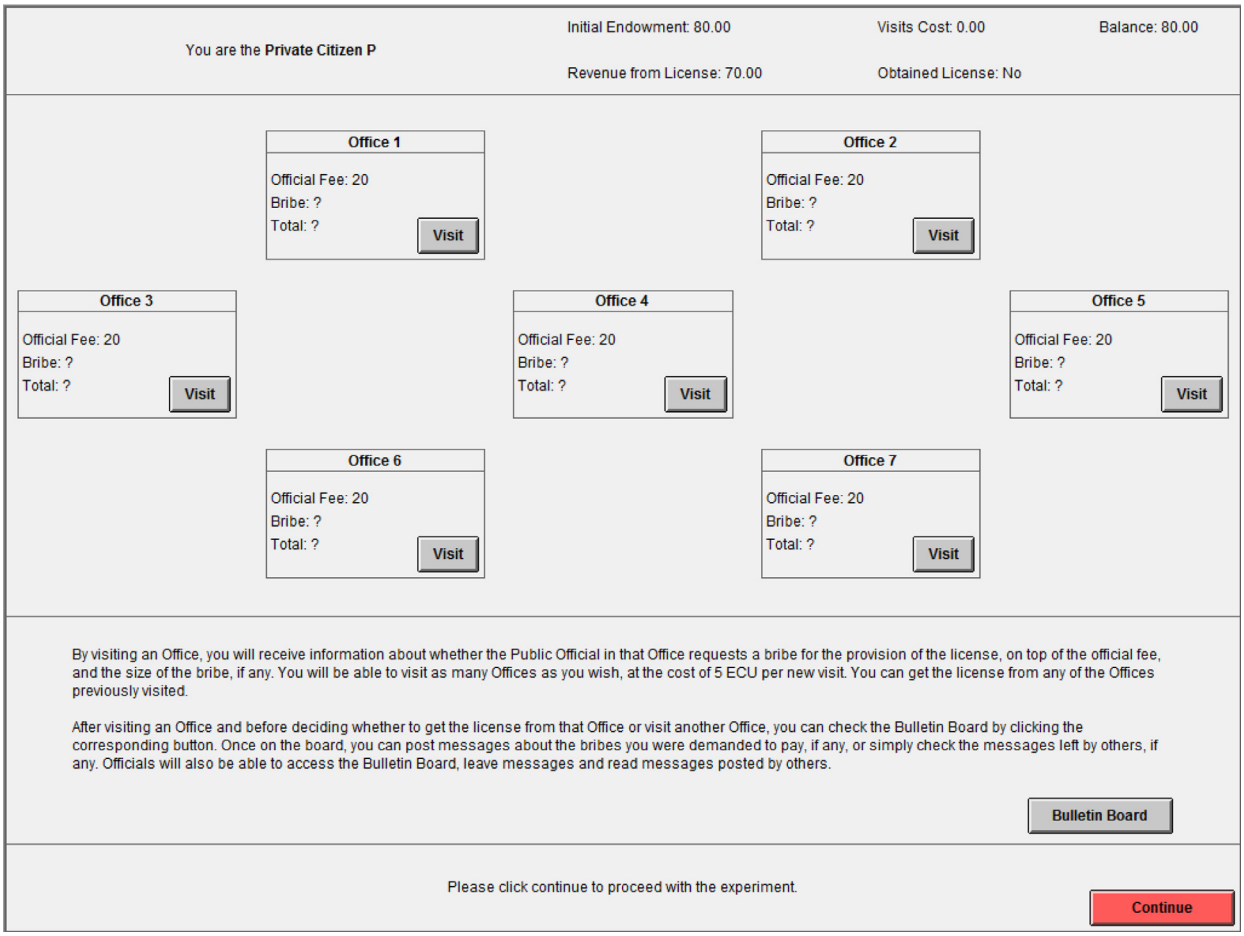

This part of the experiment is about to begin. We ask again that you not look at the screens of those around you or attempt to talk with other participants at any time during the session. You will be able to read through the instructions and click through the screens at your own pace. Each section of the experiment will begin after all participants have finished reading the instructions for that section and have clicked Continue. If you have any question about the instructions that you will receive on your screen, please feel free to raise your hand at any time during the session, and the experimenter will come to answer your questions in private.

Screen 1: Official decides whether to demand a bribe

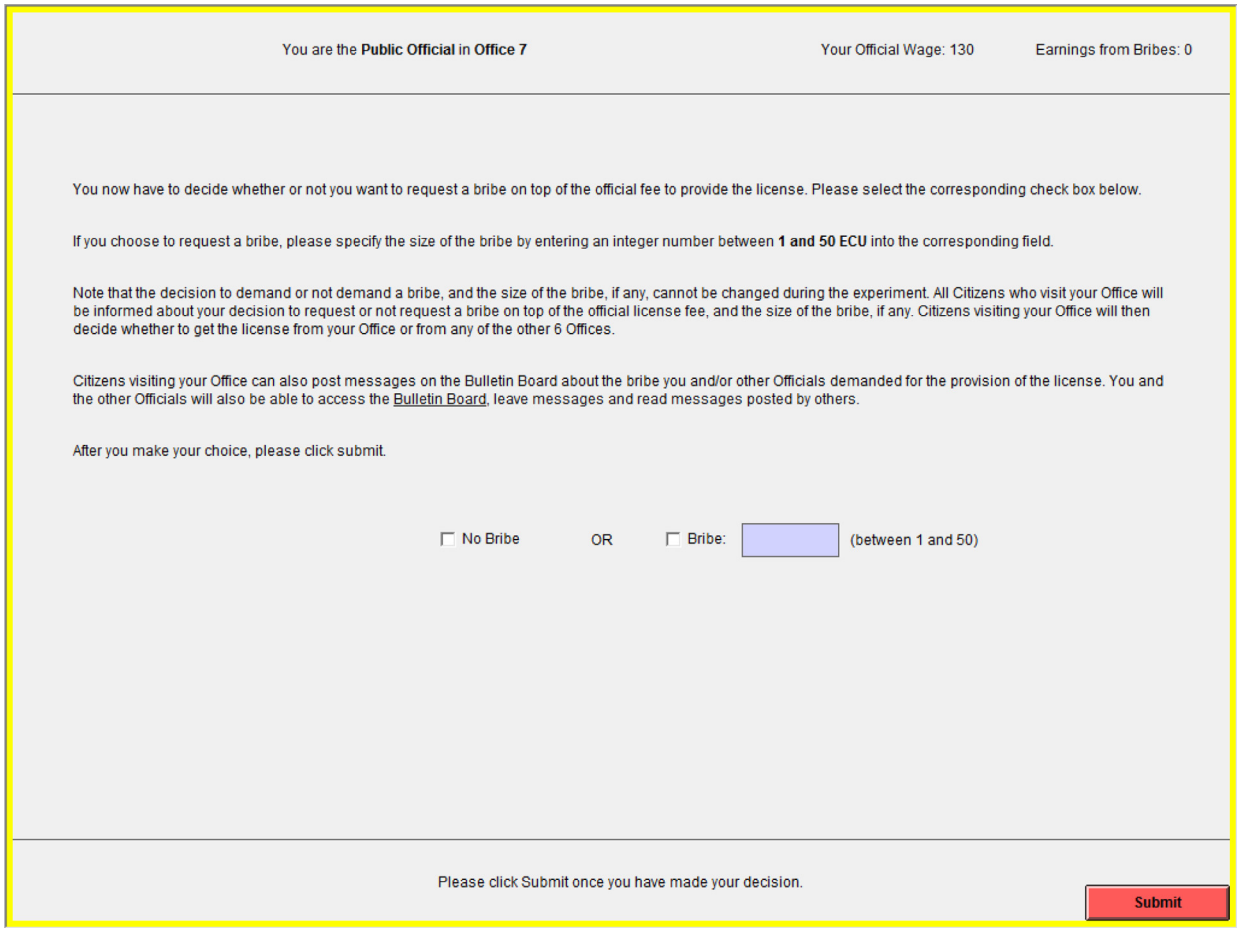

Screen 2: Citizen visits a randomly assigned office

Screen 3: Citizen decides to pay or to visit another office

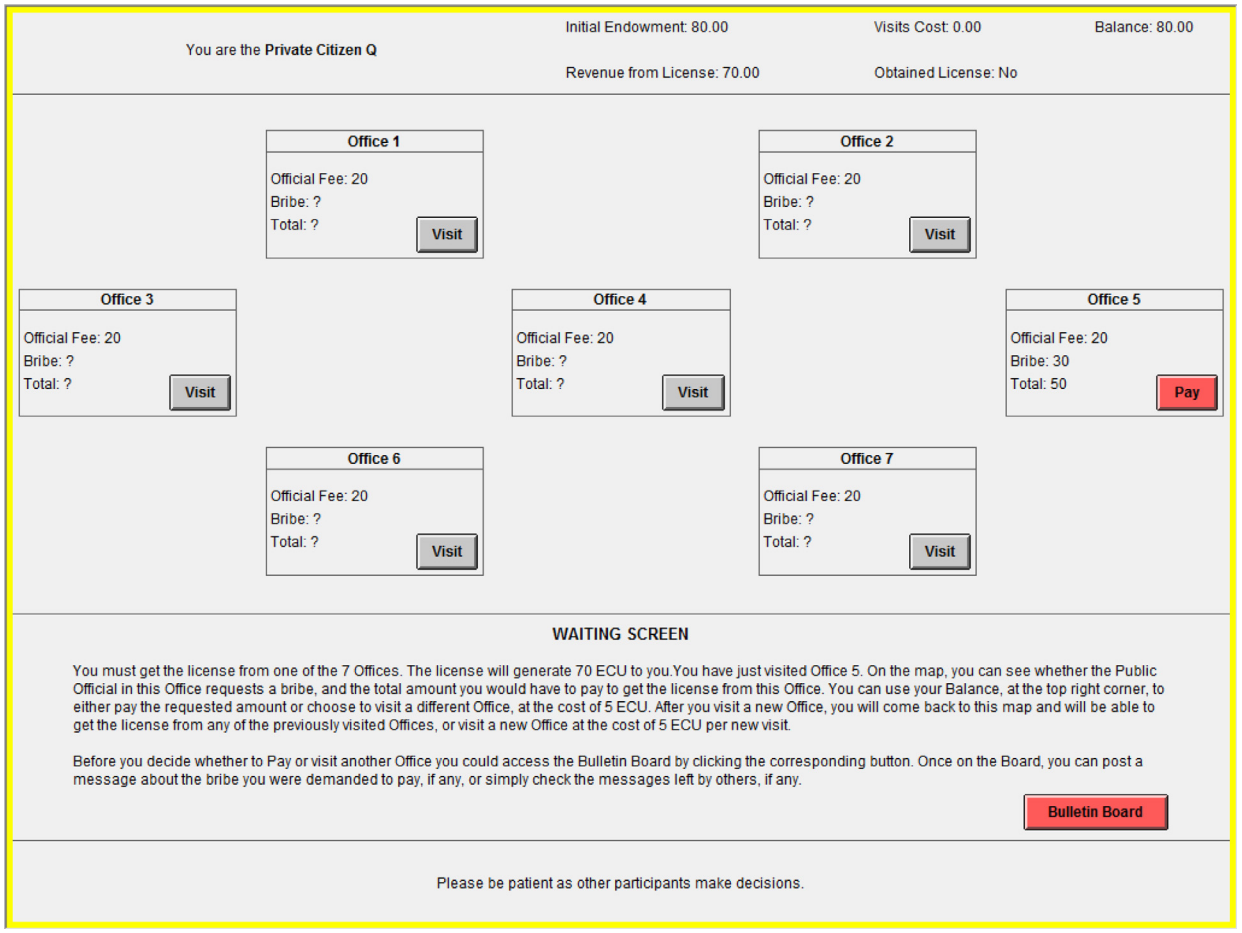

Screen 4: Official's waiting screen

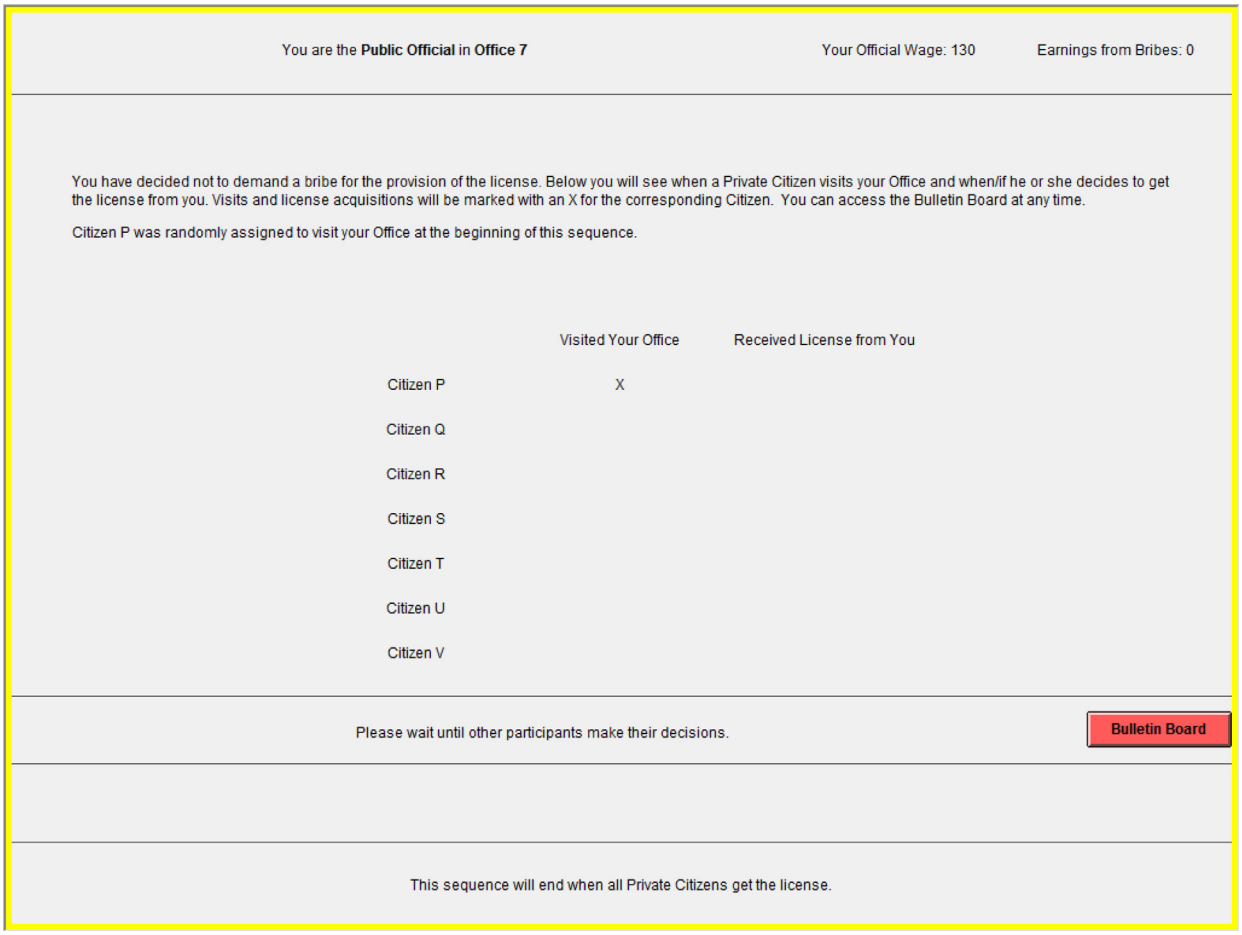

Screen 5: Bulletin board (treatments BB_CIT and BB_ALL)

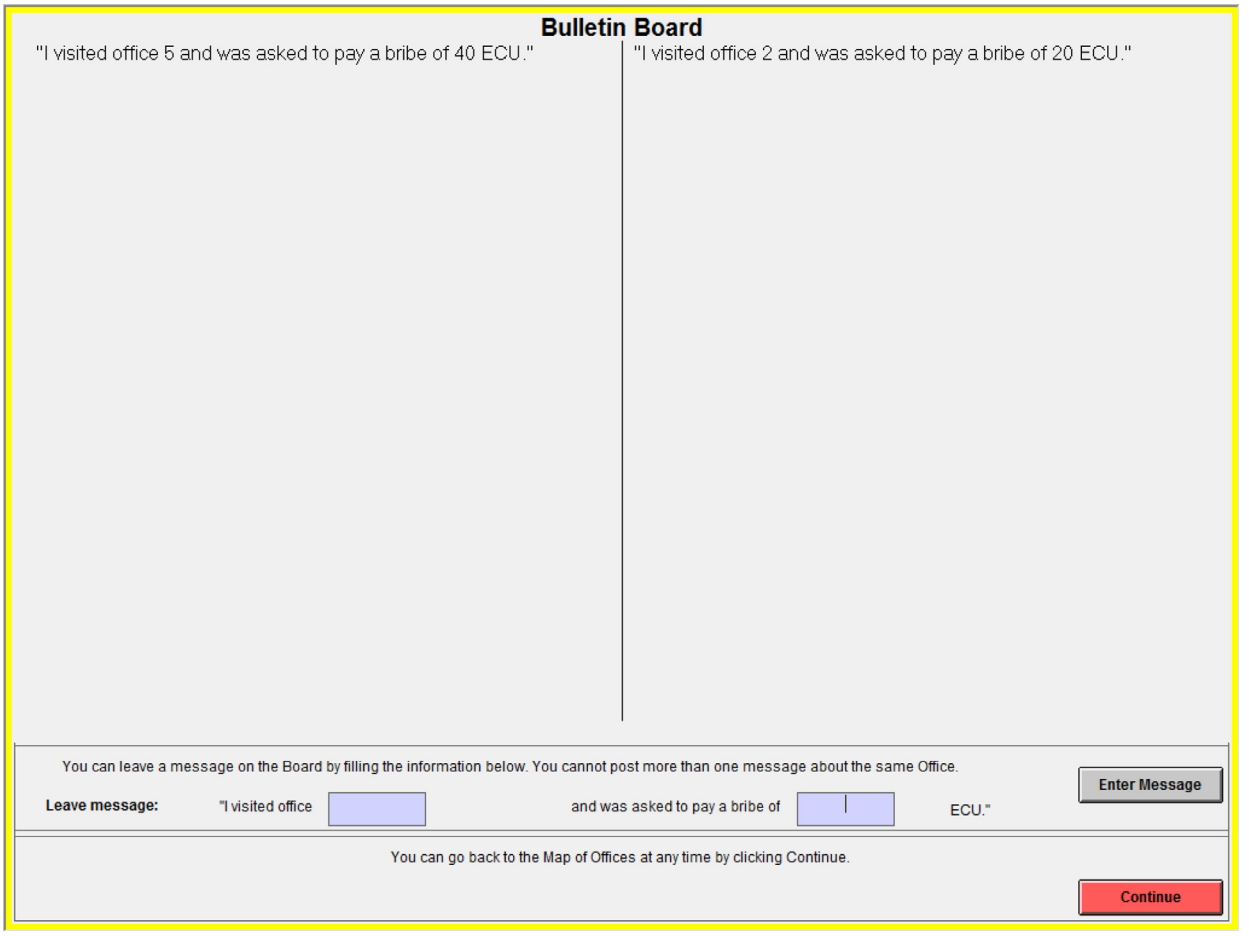

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.euroecorev.2017.02.003.

References

Abbink, K., Dasgupta, U., Gangadharan, L., Jain, T., 2014. Letting the briber go free: an experiment on mitigating harassment bribes. J. Public Econ. 111, 17–28.
Abbink, K., Irlenbusch, B., Renner, E., 2002. An experimental bribery game. J. Law Econ. Organ. 18 (2), 428–454.
Anderson, E.T., Simester, D., 2013. Advertising in a competitive market: the role of product standards, customer learning, and switching costs. J. Mark. Res. 50 (4), 489–504.
Anderson, M., Magruder, J., 2012. Learning from the crowd: regression discontinuity estimates of the effects of an online review database. Econ. J. 122 (563), 957–989.
Armantier, O., Boly, A., 2013. Comparing corruption in the laboratory and in the field in Burkina Faso and in Canada. Econ. J. 123 (573), 1168–1187.
Banerjee, A.V., Banerji, R., Duflo, E., Glennerster, R., Khemani, S., 2010. Pitfalls of participatory programs: evidence from a randomized evaluation in education in india. Am. Econ. J. Econ. Policy 2 (1), 1–30.
Banuri, S., Eckel, C., 2015. Cracking down on bribery. Soc. Choice Welf. 45 (3), 579–600.
Bardhan, P., 1997. Corruption and development: a review of issues. J. Econ. Lit. 35, 1320–1346.
Barr, A., Serra, D., 2010. Corruption and culture: an experimental analysis. J. Public Econ. 94 (11), 862–869.
Bikhchandani, S., Sharma, S., 1996. Optimal search with learning. J. Econ. Dyn. Control 20 (1), 333–359.
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 (2–3), 571–581.
Bolton, G., Greiner, B., Ockenfels, A., 2013. Engineering trust: reciprocity in the production of reputation information. Manag. Sci. 59 (2), 265–285.
Burdett, K., Judd, K.L., 1983. Equilibrium price dispersion. Econometrica 51, 955–969.
Cabral, L., Hortacsu, A., 2010. The dynamics of seller reputation: evidence from ebay. J. Ind. Econ. 58 (1), 54–78.
Camerer, C., 2011. The promise and success of lab-field generalizability in experimental economics: a critical reply to Levitt and List. Working paper available at SSRN: https://ssrn.com/abstract=1977749.
Cameron, L., Chaudhuri, A., Erkal, N., Gangadharan, L., 2009. Propensities to engage in and punish corrupt behavior: experimental evidence from Australia, India, Indonesia and Singapore. J. Public Econ. 93 (7), 843–851.
Cason, T.N., Friedman, D., 2003. Buyer search and price dispersion: a laboratory study. J. Econ. Theory 112 (2), 232–260.

Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., Rogers, F.H., 2006. Missing in action: teacher and health worker absence in developing countries. J. Econ. Perspect. 20 (1), 91–116.
Chevalier, J.A., Mayzlin, D., 2006. The effect of word of mouth on sales: online book reviews. J. Mark. Res. 43 (3), 345–354.
Davis, D.D., Holt, C.A., 1996. Consumer search costs and market performance. Econ. Inq. 34 (1), 133.
Diamond, P.A., 1971. A model of price adjustment. J. Econ. Theory 3 (2), 156–168.
Duflo, E., Dupas, P., Kremer, M., 2015. School governance, teacher incentives, and pupil–teacher ratios: experimental evidence from Kenyan primary schools. J. Public Econ. 123, 92–110.
Fischbacher, U., 2007. z-tree: Zurich toolbox for ready-made economic experiments. Exp. Econ. 10 (2), 171–178.
Gneezy, U., Rockenbach, B., Serra-Garcia, M., 2013. Measuring lying aversion. J. Econ. Behav. Organ. 93, 293–300.
Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with orsee. J. Econ. Sci. Assoc. 1 (1), 114–125.
Gupta, S., Davoodi, H., Alonso-Terme, R., 1998. Does Corruption affect Inequality and Poverty?. International Monetary Fund.
Holt, C.A., Laury, S.K., 2002. Risk aversion and incentive effects. American Economic Review 92 (5), 1644–1655.
Houser, D., Wooders, J., 2006. Reputation in auctions: theory, and evidence from ebay. J. Econ. Manag. Strategy 15 (2), 353–369.
Hunt, J., 2007. How corruption hits people when they are down. J. Dev. Econ. 84 (2), 574–589.
Kessler, J., Vesterlund, L., 2011. The External Validity of Laboratory Experiments: Qualitative Rather Than Quantitative Effects. Oxford University Press.
Lafky, J., 2014. Why do people rate? Theory and evidence on online ratings. Games Econ. Behav. 87, 554–570.
Luca, M., Zervas, G., 2013. Fake it Till You Make it: Reputation, Competition, and Yelp Review Fraud. Harvard Business School NOM Unit Working Paper 14-006.
Masclet, D., Pénard, T., 2012. Do reputation feedback systems really improve trust among anonymous traders? An experimental study. Appl. Econ. 44 (35), 4553–4573.
Mauro, P., 1995. Corruption and growth. Q. J. Econ. 110 (3), 681–712.
Mayzlin, D., Dover, Y., Chevalier, J.A., 2012. Promotional Reviews: An Empirical Investigation of Online Review Manipulation. Technical Report. National Bureau of Economic Research.
Méon, P.-G., Sekkat, K., 2005. Does corruption grease or sand the wheels of growth? Public Choice 122 (1–2), 69–97.
Morgan, J., Orzen, H., Sefton, M., 2006. An experimental study of price dispersion. Games Econ. Behav. 54 (1), 134–158.
Olken, B.A., 2006. Corruption and the costs of redistribution: micro evidence from Indonesia. J. Public Econ. 90 (4), 853–870.
Olken, B.A., 2007. Monitoring corruption: evidence from a field experiment in Indonesia. J. Political Econ. 115 (2), 200–249.
Pradhan, M., Suryadarma, D., Beatty, A., Wong, M., Gaduh, A., Alisjahbana, A., Artha, R.P., 2014. Improving educational quality through enhancing community participation: results from a randomized field experiment in Indonesia. Am. Econ. J. Appl. Econ. 6 (2), 105–126.
Reinikka, R., Svensson, J., 2004. Local capture: evidence from a central government transfer program in Uganda. Q. J. Econ. 119 (2), 679–705.
Resnick, P., Zeckhauser, R., Swanson, J., Lockwood, K., 2006. The value of reputation on ebay: a controlled experiment. Exp. Econ. 9 (2), 79–101.
Rockenbach, B., Sadrieh, A., 2012. Sharing information. J. Econ. Behav. Organ. 81 (2), 689–698.
Ryvkin, D., Serra, D., 2012. How corruptible are you? Bribery under uncertainty. J. Econ. Behav. Organ. 81 (2), 466–477.
Ryvkin, D., Serra, D., 2015. Is More Competition Always Better? An Experimental Study of Extortionary Corruption. FSU Working Paper.
Ryvkin, D., Serra, D., 2016. The Industrial Organization of Corruption: Monopoly, Competition and Collusion Among Public Officials. Mimeo.
Salmon, T., Serra, D., 2016. Corruption, social judgment and culture: an experiment. Working paper available at http://faculty.smu.edu/dserra/SalmonSerra_Culture.
Sequeira, S., 2012. Advances in Measuring Corruption in the Field. Emerald Group Publishing Limited, pp. 145–175.
Sequeira, S., Djankov, S., 2013. Corruption and firm behavior. Working paper available at http://eprints.lse.ac.uk/54321/.
Serra, D., 2006. Empirical determinants of corruption: a sensitivity analysis. Public Choice 126 (1–2), 225–256.
Serra, D., 2012. Combining top-down and bottom-up accountability: evidence from a bribery experiment. J. Law Econ. Organ. 28 (3), 569–587.
Serra, D., Wantchekon, L., 2012. New Advances in Experimental Research on Corruption, Vol. 15. Emerald Group Publishing.
Stahl, D.O., 1989. Oligopolistic pricing with sequential consumer search. Am. Econ. Rev. 79, 700–712.
Tanzi, V., Davoodi, H., 1998. Corruption, Public Investment, and Growth. Springer.
The World Bank, 2004. Making Service Work for Poor People. The World Bank, Washington, DC. World Development Report
Treisman, D., 2000. The causes of corruption: a cross-national study. J. Public Econ. 76 (3), 399–457.
Ye, Q., Law, R., Gu, B., 2009. The impact of online user reviews on hotel room sales. Int. J. Hosp. Manag. 28 (1), 180–182.