Preemptive Corruption, Hold-Up and Repeated Interactions*

Emmanuel Dechenaux† Andrew Samuel‡
Kent State University Loyola College

June 14, 2009

Abstract

This paper analyzes repeated interactions between a firm and an inspector who monitors regulatory compliance. In the model, the firm may offer a bribe to preempt the inspection. Corruption is unfeasible in the one-shot game because of inspector hold-up. In an infinitely repeated game, we characterize the set of bribes which can be sustained as equilibrium paths using a trigger strategy. Surprisingly, the most likely bribe-givers are not the firms that benefit the most from the illegal behavior. Furthermore, strengthening anti-corruption policies has ambiguous welfare effects because it only improves compliance among a subset of firms and increases monitoring effort.

Keywords: Corruption, Collusion, Bribery, Law enforcement, Repeated game
J.E.L. classification: D21, K42

*We wish to thank Nicola Lacetera, Justin Sydnor, Joel Watson as well as seminar participants at Case Western Reserve University, the Annual Meetings of the Midwest Economics Association in Cleveland (March 2009) and the American Law and Economics Association in San Diego (May 2009) for helpful comments.

†Corresponding author. Department of Economics, Kent State University, Kent, OH 44242. Phone: 330-672-1106. Fax: 330-672-9808. Email: edechena@kent.edu.

‡Department of Economics, Loyola College in Maryland, 4501 N. Charles Street, Baltimore, MD 21210. Phone: 410-617-5328. Email: asamuel@loyola.edu.

1 Introduction

This paper analyzes corruption between an inspector and an agent in a situation where the inspector must exert costly effort to imperfectly monitor the agent. Collusion in the form of bribery, referred to as corruption, may occur before the inspector spends costly monitoring effort (preemptively) or after the inspector incurs the cost (ex-post). For example, consider the case of corruption between border guards and opium smugglers in Afghanistan (Lubin, 2003). Lubin describes scenarios where drug smugglers choose to either preemptively bribe the guards or to do so ex-post only if they are caught smuggling opium. A one-shot game modeling this situation is analyzed by Samuel (2008) who builds on a model first developed by Mookherjee and Png (1995). In this model, the agent is a firm that must incur a cost of compliance and thus, has an incentive to bribe the inspector in charge of monitoring compliance.[1] Mookherjee and Png (1995) study the welfare implications of ex-post bribery and examine the effectiveness at curbing corruption of policy parameters such as the inspector's reward for producing evidence of illegal behavior by the firm. Samuel (2008) revisits these issues but allows for preemptive as well as ex-post corruption (see also Motta, 2009). However, there is a key issue that neither paper addresses. Indeed, corrupt contracts are not enforceable in a court of law and may therefore suffer from the hold-up problem. That is, an inspector who accepted a bribe in exchange for not inspecting or reporting a firm may nonetheless report or fine the firm after receiving the bribe. We therefore believe that repeated play is important in this context because it helps solve the hold-up problem that exists both with preemptive as well as ex-post collusion.

Although there is a vast literature on the type of bureaucratic or administrative corruption we analyze, few papers have examined the problem of hold-up and repeated interactions (Bardhan, 2005). Most of the theoretical research on corruption has focused on one-shot interactions. A brief summary of the key results from this literature is as follows. Becker and Stigler (1974) argue that the privatization of law enforcement eliminates bribery, but Mookherjee and Png (1995) show that it is not always socially optimal because it encourages inspectors to exert too much effort. Furthermore, as long as inspectors and agents are homogeneous, it is always socially optimal to eliminate ex-post bribery completely (Mookherjee

[1] In the remainder of the paper, we interchangeably refer to lack of compliance by the firm as "choosing the bad technology" or "engaging in illegal behavior."

and Png 1995; Besley and MacLaren, 1993).[2] Finally, when bribery can occur either preemptively or ex-post, then privatizing law enforcement does not necessarily eliminate bribery and, more importantly, it is not always socially optimal to eliminate preemptive bribery, while it is always optimal to eliminate ex-post bribery (Samuel, 2008).

A shortcoming of the previous literature is the implicit assumption that the collusive contract can be enforced. Most collusive contracts, however, are not legally enforceable, and therefore are very likely to suffer from the hold-up problem as recognized in a number of recent studies (Jain, 2001; Lambsdorff, 2002; Klochko and Ordeshook, 2003, and Lambdsorff and Nell, 2007). When hold-up is possible, then unless the inspector or the firm have some credible threat to punish deviations from the corrupt agreement, bribery is not feasible. In this context, repeated interactions seem a natural candidate to consider as an enforcement device. For instance, Mookherjee and Png (1995) write:

> "*(t)here are several reasons to think that the illicit contract between inspector and factory can be enforced. First, there may be a continuing relationship between the inspector and the factory, so that long-term costs of reneging outweigh the one-shot gain,*" (page 150)

while Lambsdorff (2002) devotes an entire section of his survey to the importance of repeated interactions.

Leniency policies for whistleblowers may also serve to enforce bribe contracts that would otherwise not be feasible because of hold-up (Buccirossi and Spagnolo, 2006; Lambsdorff and Nell, 2007).[3] However, while it is clear that leniency policies can exacerbate the occurrence of corruption, they cannot possibly be the root cause of the problem they were designed to alleviate. Therefore, in our model, we do not consider leniency policies and consequently, bribery does not occur in the one-shot game (see also Lambert-Mogiliansky et al., 2007 and Klochko and Ordeshook, 2003).

In the repeated game, we find that preemptive bribery does occur because players can use a trigger strategy to punish deviations from the collusive path. This result is interesting in light of a recent study on bribery in the trucking industry in the Indian state of Gujurat

[2] In a similar model, Besley and MacLaren (1993) also show that it is not optimal to eliminate bribery completely when agents or inspectors are heterogeneous.

[3] Motta and Polo (2003) analyze the effect of leniency policies on cartel behavior.

(Center for Electronic Governance, 2002). Trucks may either be over or under the axle-load limit and the department of transportation authorizes weigh-station inspectors to inspect and fine trucks that violate the load limit. In order to avoid these fines, most truckers pay bribes before their trucks are inspected, i.e. preemptively. Indeed, the same study also points out that most of the trucks examined in the survey were those that regularly "*cross and re-cross the check-posts in Gujurat*" (Page 5). Thus, it is very likely that repeated interactions do indeed play a role in enforcing the preemptive bribe contract.[4]

Interestingly, we find that the firms that are most likely to offer preemptive bribes are not those which benefit the most from engaging in the illegal activity, nor those which benefit the least. Thus, in contrast with static models where the firm's profit from engaging in the illegal activity has no effect on the incentives for corruption, in our dynamic model, the firm's profit matters because it is lost whenever an audit finds evidence of corruption. This result is particularly interesting because it is related to case study evidence and an empirical finding that different types of firms make different choices with regard to corruption (Lubin, 2003; Svensson, 2003). Svensson (2003) finds that the firms that are most able to pay, as measured by expected profitability, pay higher bribes. Additionally, in his study, firms with greater refusal power, as measured by a low cost of having to exit the market, are less likely to pay bribes. In our model, conditional on paying a bribe, firms that benefit more from the illegal activity are willing to pay more. Furthermore, firms that benefit very little from the illegal activity have great refusal power since their cost of full compliance is low. It may be argued that the firms which earn the greatest benefit from the illegal activity also have a high degree of refusal power because, even when fined by the regulator, their expected payoff is high. We show that firms at these two extremes of the private benefit distribution are the least likely to pay preemptive bribes.[5]

[4]Because our results build on Mookherjee and Png (1995), whose model is based on Becker and Stigler (1974), they also apply to the situation studied by Di Tella and Schargrodsky (2003). In their paper, a hospital purchase manager may accept a bribe from a specific supplier. In exchange, the purchase manager foregoes search for lower price quotes. With minor changes, our model fits the situation with the inspector in the role of the purchase manager and the firm in the role of the high-price supplier.

[5]However, one must be cautious in interpreting these findings in light of our results. In our model, there is imperfect monitoring, but the legal environment is one in which the regulator can credibly commit to enforce the law conditional on obtaining appropriate information. That is, if evidence is found and reported, then the firm is fined with probability one. Furthermore, if corruption is detected, then penalties are applied with probability one. It is not clear that these assumptions hold for the firms in Svensson's (2003) dataset. Furthermore, for the firms in Svensson's analysis, it is is not possible to tell whether bribery was the result

Our analysis offers a number of comparative static results that may be useful in assessing the value of anti-corruption policy instruments such as the penalty for bribery and the supervisor's reward for monitoring. Contrary to the previous literature, our analysis of the feasibility of preemptive corruption reveals that an increase in the probability of detecting corruption and a decrease in the penalties for corruption that leave the expected fines for bribery unchanged unambiguously reduces the incentives for corruption. This result does not arise in existing static models and suggests that the effects of anti-corruption policies when bribery is an equilibrium in a repeated game are distinct from those in a static game.

Moreover, the welfare effects of anti-corruption policies are not straightforward to characterize because of our result that firms that engage in preemptive bribery are not those that benefit the most, nor those that benefit the least from the illegal activity. Rather, the firms that pay preemptive bribes are in an intermediate range of the private benefit distribution. As penalties for corruption are raised, fewer firms engage in preemptive bribery and this range shrinks. Some of the firms that previously engaged in preemptive bribery continue to pursue the illegal activity facing the risk of detection by the inspector, while other firms switch from corruption to compliance. For the former firms, increasing penalties for corruption generates higher monitoring cost without lowering the harm associated with the illegal activity, thereby decreasing overall welfare. For the latter, strengthening anti-corruption policies reduces both bribery and the illegal activity, thereby increasing social welfare. Because of this heterogeneity in responses to changes in anti-corruption policies, the overall welfare effects are not as straightforward to characterize as in the one-shot game.

Finally, rather than fully analyze ex-post corruption, we show that if the firm and the inspector can coordinate using a public randomization device whereby, with some probability a preemptive bribe is exchanged and otherwise no bribery or inspection occurs, then ex-post bribery will never occur.[6] In particular, we show that the incentive for hold-up on the ex-post collusive path is no less than the incentive to hold up on the randomized bribery path with public randomization. More importantly, we also show that for every feasible preemptive bribe scheme, there exists a scheme with a public randomization device that

of capture or extortion. As we mention below, we focus on capture and do not analyze extortion.

[6]The report on bribery in the trucking industry in Gujurat mentioned above interestingly notes that weigh-station inspectors often randomly allowed trucks to bypass the weigh-station in exchange for an unofficial payment or bribe (Center for Electronic Governance, 2002).

Pareto dominates preemptive bribery without randomization. This result is noteworthy given that most of the literature has focused exclusively on ex-post bribery and collusion (for example, Mookherjee and Png, 1995; Polinsky and Shavell, 2003; Acemoglu and Verdier, 2000; Strausz, 1999). Perhaps that is because intuition suggests preemptive corruption is more susceptible to the hold-up problem and thus is unlikely to occur. However, our model shows that even when hold-up is a concern, preemptive bribery can be sustained and, moreover, it will always be chosen over ex-post bribery if the firm and the inspector have access to a public randomization device.

In the existing literature, Choi and Thum (2004) and Lambert-Mogiliansky et al. (2007) also consider a version of the hold-up problem. In these studies, hold-up refers to the behavior of bureaucrats who refuse to officially approve otherwise qualified business projects unless they are paid a bribe. Thus, hold-up is akin to the extortion of honest entrepreneurs. On the other hand, in our model, bureaucrats are law-enforcers who always have the option to report a dishonest or unqualified agent, thereby reneging on their promise to turn a blind eye on wrong-doing by the agent. Lambert-Mogiliansky et al. (2007) refer to the practice of demanding a bribe to approve unqualified projects as capture. To understand where our analysis fits within this framework, note that in our paper corruption occurs in the form of capture, while we refer to hold-up as the practice of reneging on the promise made in any corrupt contract, preemptive or ex-post.

The paper is organized as follows. Section 2 outlines the one-shot extensive form game, which we refer to as the Inspection game. Section 3 presents the infinitely repeated game and derives results regarding the sustainability of preemptive corruption. Section 4 discusses the welfare effects of anti-corruption policies in the context of our model. Section 5 analyzes preemptive corruption with a public randomization device and Section 6 concludes. All proofs appear in the Appendix.

2 The One-Shot Game

In this section, we describe the simple extensive form game, a closely related version of which has been analyzed by Samuel (2008). His model is itself based on Mookherjee and Png (1995). Formally, there are three risk-neutral actors: the principal or the regulator, the

inspector (S), and the firm (F). The regulator does not make any strategic decisions and solely rewards the inspector and applies penalties in accordance with the law. The extensive form is given in Figure 1. In the continuation, we refer to this extensive form game as the *Inspection game*. Pure strategies are easily defined and we denote them by σ_F for the firm and σ_S for the inspector.[7] The players' expected payoffs in the Inspection game are denoted by $U_F(\sigma_F, \sigma_S)$ and $U_S(\sigma_F, \sigma_S)$.

The firm determines its level of emissions through its choice between $W = 0$, the good technology, and $W = 1$, the bad technology. Regulation stipulates that the firm must use the good technology. If the firm chooses $W = 1$, it obtains a private benefit of $w > 0$ while a choice of $W = 0$ yields no private benefit. The social harm or external cost from choosing $W = 1$ is equal to h, where $h > 0$. Since the firm maximizes its private benefit, absent regulation, it would always choose the bad technology. With imperfect information regarding the firm's choice of W and costly monitoring, the regulator hires inspectors to ensure that firms comply with the regulation.[8] If the firm is found guilty of not complying, it is required to pay a fine f, which is set by the regulator.[9] It is implicitly assumed that the inspector cannot plant evidence. That is, if the firm chooses $W = 0$, it does not have to offer a bribe and does not pay any fine.

As in Samuel (2008), upon visiting the firm, the inspector immediately observes whether the firm has chosen the good or the bad technology. This knowledge of W is soft, and not verifiable by a third party. Therefore, in order to fine the firm, he must obtain hard, verifiable evidence by exerting effort. In this model, the inspector is not paid directly by the regulator but is given a payment of rf; that is, he is given a fraction r of the fines he is allowed to collect ($r \leq 1$). In order to monitor the firm with intensity μ the inspector must exert an effort level of $e(\mu)$. The intensity $\mu \in [0, 1]$ represents the probability that the inspector will be able to discover and produce verifiable evidence on whether the firm's

[7] Let $\mathcal{I}_i$ be the set of information sets for player i, $i \in \{F, S\}$ with generic element I_i and $D_i(I_i)$ be the set of actions available to player i at information set I_i. Then $D_i = \cup_{I_i \in \mathcal{I}_i} D(I_i)$ is the set of actions available to player i. A pure strategy σ_i for player i is a function $\sigma_i : \mathcal{I}_i \to D_i$ with $\sigma_i(I_i) \in D(I_i)$ for all $I_i \in \mathcal{I}_i$. Furthermore, in the remainder of the paper, we will use the following notation for actions available to the players. The variable B denotes a preemptive bribe and b denotes an ex-post bribe. Following a preemptive bribe, we let $A \in \{\text{Accept}, \text{Reject}\}$. Similarly, following an ex-post bribe offer, we let $a \in \{AR, AD, R\}$, where $AR =$ "Accept and Report", $AR =$ "Accept and Destroy" and $R =$ "Reject".

[8] Since the principal is the regulator, she can credibly commit to monitor the firm ex-ante.

[9] The interested reader is referred to Samuel (2008) for a description of how these assumptions differ from those in Mookherjee and Png (1995).

choice of W is equal to 1 or if instead $W = 0$. If the inspector discovers that $W = 1$, he is authorized to fine the firm and receives a reward of rf. We make the following assumptions on $e(\mu)$.

Assumption 1 *The cost function $e(\mu)$ is continuous, strictly increasing, strictly convex and differentiable twice. Furthermore, $\lim_{\mu\uparrow 1} e'(\mu) = \infty$ and $e(0) = e'(0) = 0$.*

The inspector is corruptible, so he may accept a bribe from the firm. In any given interaction between a firm and an inspector there are two stages in which collusion may occur. In the first stage, the firm can offer a preemptive bribe, denoted by B, before the inspector has exerted any effort. Alternatively, collusion may occur in the second stage. That is, the inspector first exerts effort and then, if he finds evidence of the firm's non-compliance, the firm may offer an *ex-post* bribe. In the remainder of the paper, we refer to the exchange of any strictly positive bribe as corruption. In both types of corruption, upon receiving the bribe, the inspector promises to report that the firm's technology is $W = 0$ and does not fine the firm at all. In the event that either one of the two bribes is exchanged, with probability λ an audit by the regulator may find evidence of corruption, in which case the regulator penalizes both parties for their wrongful transaction.[10] A firm (the bribe giver) found engaging in collusion in either or both stages is fined a finite amount $p_g f$ in addition to having to pay the fine f. Similarly, the inspector (the bribe taker) is penalized with a finite fine p_t. For simplicity, we assume that the penalties do not depend on the size of the bribes and thus, are the same whether or not both types of corruption occurred.[11]

In order to establish useful notation, consider the optimal choice of monitoring intensity by an inspector who found that the firm's technology is $W = 1$, but anticipates that no ex-post bribe will be exchanged. In this case, the inspector's expected payoff is equal to $\mu rf - e(\mu)$. Under Assumption 1 and given $rf > 0$, the optimal choice of monitoring intensity is μ_n such that $rf = e'(\mu_n)$. The optimal value of μ is unique and satisfies $0 < \mu_n < 1$.

[10] Here, λ is fixed, like in Becker and Stigler (1974) and in most papers that build on their model, including Samuel (2008). See Besley and McLaren (1993) for a short discussion of the effect of endogenizing the probability of detection in a similar model.

[11] Clearly, in both Mookherjee and Png (1995) and our model, all corruption could be eliminated by setting the penalties sufficiently high. However, this may imply setting fees at arbitrarily high levels, which would likely be unfeasible in reality. Furthermore, because these fines may have to fit any situation involving corruption, they are generic and cannot depend on w or h. These assumptions are similar to Mookherjee and Png (1995).

Furthermore, define $e(\mu_n) \equiv e_n$. Proposition 1 below characterizes equilibrium behavior and establishes the straightforward result that due to hold-up, corruption will not occur in the subgame perfect Nash equilibrium (SPNE).[12]

Proposition 1 *The Inspection game has a unique SPNE and corruption never occurs in equilibrium.*

If $w - \mu_n f \geq 0$, then on the equilibrium path, the firm chooses $W = 1$, the inspector exerts effort e_n, and fines the firm an amount f with probability μ_n. In equilibrium, the firm's expected payoff is equal to $U_F^N = w - \mu_n f$ and the inspector's expected payoff is equal to $U_S^N = \mu_n r f - e_n$.

If $w - \mu_n f < 0$, then on the equilibrium path, the firm chooses $W = 0$ and both the firm and the inspector earn a payoff of zero.

Proof. See the Appendix. ■

The result in Proposition 1 is in sharp contrast to the equilibrium result of Mookherjee and Png (1995) and Samuel (2008) because both assume that hold-up does not arise. Proposition 1 establishes that if the inspector and the firm only interact once, then bribery will not occur. The goal of this paper is to model a plausible mechanism by which hold-up is prevented and bribery can be supported as an equilibrium. The analysis below shows that explicitly modeling how repeated interactions are used to sustain collusive equilibria yields comparative statics result that differ from those obtained in the static models. Before we proceed with the analysis, we state the following assumption.

Assumption 2 *The firm's private benefit w satisfies $w \geq \mu_n f$.*

According to Proposition 1, Assumption 2 implies that in the SPNE of the component game, the firm chooses the bad technology, the inspector investigates the firm's behavior and imposes a fine of f with probability μ_n. Furthermore, the respective players earns payoffs equal to $U_F^N = w - \mu_n f$ and $U_S^N = \mu_n r f - e_n$. Unless otherwise mentioned, we maintain Assumption 2 for the remainder of the analysis.

[12]Since we focus on pure strategies, if the firm is indifferent between $W = 0$ and $W = 1$, we assume it chooses $W = 1$.

3 An infinite-horizon model of preemptive corruption

We now consider the possibility of repeated interactions between the firm and the inspector. The model is a simple infinite-horizon game with complete information and discounting, in which the Inspection game analyzed in the previous section is repeated unless the regulator decides to put an end to the interaction between the two players. Below, we formally describe the game.

3.1 The game and definitions

Suppose now that there are a countably infinite number of periods indexed by t and that in every period, the firm must chooses its technology ($W = 1$ or $W = 0$). Following the firm's choice, the inspector investigates the firm's actions like in the Inspection game of the previous section. Importantly, we assume that an audit which reveals that bribery occurred puts an end to the interaction between the firm and the inspector. In this case, both the firm and inspector earn a per-period payoff of zero in the continuation of the game.[13] Like in the component Inspection game, if the firm and the inspector exchanged a bribe, then they are caught with probability λ as a result of the audit. Importantly, we assume that bribery in period t cannot be detected in period $s > t$. That is, the evidence of corruption vanishes at the end of period t.

In this game, a path τ is an infinite sequence of actions $\{(W_s, B_s, A_s, \mu_s, b_s, a_s)\}_{s=0}^{\infty}$. To keep things simple, we assume that for both the firm and the inspector, the observable history at the beginning of period $t > 0$ consists of the sequence $\{(W_s, B_s, A_s, \mu_s, b_s, a_s)\}_{s=0}^{t-1}$. That is, the firm's choice of W and the inspector's choice of μ are observable both by the inspector and the firm. However, we continue to assume that the regulator is unable to observe μ and thus cannot condition the inspector's tenure on his choice of μ. Therefore, in every period t in which the firm and the inspector interact, the regulator observes any bribe that is exchanged with probability λ and terminates the relationship between the firm and the inspector if either B_t, or b_t, or both are strictly positive. We call a path *stationary* if the players use the same actions in every Inspection game on the path, that is

[13] We are assuming that if the regulator finds evidence of bribery, the inspector is dismissed and earns a reservation wage of zero, while the firm is so tightly monitored that it is forced to choose $W = 0$ in all subsequent periods.

$(W_s, B_s, A_s, \mu_s, b_s, a_s) = (W, B, A, \mu, b, a)$ for every s. Furthermore, there are paths on which some information sets are not reached. When this is so, for simplicity, we omit the actions associated with information sets that are not reached from the definition of the path.

Let $U_i(t, \tau)$ be player i's expected payoff in the Inspection game in period t of the path τ, for i in $\{F, S\}$. When there is no ambiguity, we use the abbreviation $U_i(t)$ to denote $U_i(t, \tau)$. We assume that the firm and the inspector discount future payoffs using a common discount factor δ satisfying $0 < \delta < 1$.[14] Player i's discounted cumulative expected payoff in period t of path τ in the infinitely repeated game is thus

$$V_i(t, \tau) = U_i(t) + \sum_{s=1}^{\infty} \delta^s \prod_{k=t}^{t+s-1} (1 - P(k)) U_i(t + s)$$

where the probability of termination in period k, $P(k)$, is given by

$$P(k) = \begin{cases} \lambda & \text{if } \max\{B_k, b_k\} > 0, \\ 0 & \text{if } B_k = b_k = 0. \end{cases}$$

In the special case where τ is a stationary path that involves a bribe, $P(k) = \lambda$ for every k and $U_i(t) = U_i$ for every t. Thus individual i's expected payoff can be simplified as follows

$$V_i(t, \tau) = \sum_{s=0}^{\infty} \delta^s (1 - \lambda)^s U_i = \frac{U_i}{1 - \delta(1 - \lambda)}.$$

On the other hand, if τ is stationary, but does not involve a bribe, individual i's expected payoff is:

$$V_i(t, \tau) = \sum_{s=0}^{\infty} \delta^s U_i = \frac{U_i}{1 - \delta}.$$

Since our analysis focuses on stationary paths, we make extensive use of the expected payoff functions defined by the last two equations above.

Definition 1 (Collusive Corruption) *A path τ^c is a collusive path if each player receives a payoff on τ^c that is at least as great as the discounted sum of its Inspection game equilibrium payoff and at least one player receives a strictly greater payoff.*

[14] This discount factor may also include the probability that the game continues each period. Thus, the game does not necessarily require that both players live for infinite periods.

In the following section, we analyze the feasibility of preemptive corruption, whereby the firm chooses $W = 1$, the inspector and the firm exchange a collusive preemptive bribe B, but the inspector spends no effort on inspection and reports that the firm used the good technology. We will focus on stationary paths that are collusive in the sense of Definition 1.

Since corrupt behavior is not subgame perfect in the Inspection game, to be part of a SPNE of the repeated game, collusive paths must be supported by punishment strategies that make them immune to deviations. The analysis of repeated extensive form games can be challenging because of the richness of the strategic environment (Mailath and Samuelson 2006, Section 5.4). In particular, constructing optimal punishment strategies may be an arduous task. Therefore, in this paper we focus on the following tractable punishment strategies. Let τ be the initial path. Suppose player i deviates from τ at some information set reached in the period-t Inspection game on τ. Then we assume that both players behave according to sequential rationality in the continuation of the period-t Inspection game (that is, they use actions that are optimal at each of their information sets) and they revert to the one-shot SPNE from period $t + 1$ on. Thus, a unilateral deviation by player i early on in the period t Inspection game triggers a series of (optimal) deviations in the continuation of this game. From period $t + 1$ on, the players employ the extensive-form counterpart of the grim-trigger strategy. Below we formally define such a strategy, which we refer to as *Extensive Trigger Strategy.*

Definition 2 (Extensive Trigger Strategy) *Suppose the initial path is τ and for $i \in \{F, S\}$, let $I_{i,t}$ be an information set for player i that is reached in the period-t Inspection game on path τ. For each i, the extensive trigger strategy in the infinite-horizon game specifies: Play the actions prescribed by τ as long as no deviation occurs. If a deviation occurs at information set $I_{j,t}$, $j \in \{F, S\}$, play the subgame perfect equilibrium of the subgame starting from the information set the deviation lead to. Then, conditional on not being caught by the regulator in period t, from period $t + 1$ on, play the SPNE of the Inspection game.*

It is clear that conditional on not being caught by the regulator in period t, following a deviation in period t, the firm and inspector's discounted expected payoffs from $t+1$ on are equal to $\frac{\delta}{1-\delta}U_F^N$ and $\frac{\delta}{1-\delta}U_S^N$, respectively. Also, it is clear that the punishment phase of the strategy is subgame perfect. Letting U_F^* and U_S^* denote expected payoffs from an optimal

one-period deviation from a stationary path τ and focusing on extensive trigger strategies as defined above, τ is a *stationary perfect equilibrium path* if the following two conditions hold.

$$V_S(\tau) \geq U_S^* + [1 - P(t)]\frac{\delta}{1-\delta}U_S^N, \tag{1}$$

$$V_F(\tau) \geq U_F^* + [1 - P(t)]\frac{\delta}{1-\delta}U_F^N. \tag{2}$$

3.2 Preemptive corruption equilibria

We focus on stationary collusive paths τ^c on which the firm chooses to pollute in every period ($W = 1$) and offers to pay a preemptive bribe $B^c > 0$. Such a path is a stationary perfect equilibrium path of the repeated game if and only both (1) and (2) are satisfied. We will see that this requires that the inspector has no incentive to hold the firm up and that the firm has no incentive to either choose $W = 0$ or $W = 1$ and $B < B^c$.

According to Proposition 1, if $w \geq \mu_n f$, in the SPNE of the component game, the firm chooses the bad technology, the inspector inspects the firm and imposes a fine equal to f with probability μ_n. Furthermore, the respective players earn payoffs equal to $U_F^N = w - \mu_n f$ and $U_S^N = \mu_n r f - e_n$. Otherwise, the firm chooses $W = 0$ and both players earn a payoff of zero.

To characterize collusive perfect equilibrium preemptive bribes, it is necessary to determine the players' payoffs from an optimal deviation from the path τ^c. To illustrate the incentives involved, suppose Assumption 2 holds. First, consider the firm. It is straightforward to show that it cannot earn more than U_F^N by deviating from the path. Indeed, if it chooses $W = 0$, then its payoff is equal to zero, which is less than or equal to U_F^N. On the other hand, if it offers $B < B^c$, then given the type of punishment strategies employed, the inspector will optimally accept the bribe, but spend effort e_n. Hence, the firm obtains a payoff that is less than or equal to U_F^N from the deviation. Second, consider the inspector. An optimal deviation for the inspector consists in holding the firm up by accepting B^c, but spending some effort to find hard evidence against the firm in the hope of collecting a reward. Clearly, it is optimal for the inspector to choose monitoring intensity μ_n to earn $\mu_n r f - e_n$ in the continuation of this period's game. Therefore, the inspector's payoff from an optimal deviation in period t is equal to $B^c + \mu_n r f - e_n - \lambda p_t$.

Proposition 2 below provides a necessary and sufficient condition for a preemptive bribe to be sustainable as a perfect equilibrium. Depending on the size of w (Assumption 2), if either equation (3) or (4) below is satisfied, then there exists a range of bribes that are collusive and can be sustained as a subgame perfect equilibrium using the extensive trigger strategy to punish deviations.

Proposition 2 *Suppose Assumption 2 holds. Then, there exists a collusive stationary perfect equilibrium path* $\tau^c = \{(1, B^c, 0, 0)\}$ *if and only if*

$$w - \lambda(1 + p_g)f - \lambda p_t \geq \left(\frac{1 - (1 - \lambda)\delta}{1 - \delta}\right) U_F^N + G(\lambda, \delta)U_S^N, \tag{3}$$

where $G(\lambda, \delta) \equiv \frac{1-\delta+\lambda\delta^2-\lambda^2\delta^2}{\delta(1-\delta)(1-\lambda)}$.

Suppose Assumption 2 does not hold. Then there exists a collusive stationary perfect equilibrium path $\tau^c = \{(1, B^c, 0, 0)\}$ *if and only if*

$$w - \lambda(1 + p_g)f - \lambda p_t \geq \frac{1 - \delta(1 - \lambda)}{\delta(1 - \lambda)}(\mu_n rf - e_n). \tag{4}$$

Proof. See the Appendix. ■

The left-hand sides of both (3) and (4) represent the per-period total surplus from corruption, which is equal to the benefit from engaging in the illegal activity minus the expected penalties. The right-hand side of (3) depends on U_F^N, the firm's one-period payoff from refusing to collude, as well as U_S^N, the inspector's one-period deviation payoff. The right-hand side of (4) depends solely on $\mu_n rf - e_n$, the inspector's payoff from an optimal deviation, because the one-period SPNE payoffs are equal to zero when Assumption 2 is not satisfied.

At this point, it is important to determine ranges of the parameters for which equation (3) or equation (4) hold so that preemptive corruption is a collusive equilibrium. We focus on the interesting case where Assumption 2 holds and determine ranges of the discount factor δ for which equation (3) is satisfied. Figure 2 shows three examples of the dependence of equation (3) on δ. In the figures, R stands for the right-hand side of equation (3), while L is the left-hand side. Note that L does not depend on δ. The three figures correspond to three different values of λ. In Figure 2a, λ is strictly positive but relatively small. In this case,

collusion is feasible for values of δ between $\tilde{\delta}_1$ and $\tilde{\delta}_2$. These critical values of δ can be shown to be bounded away from 0 and 1 respectively whenever $\lambda > 0$. Therefore, we show that when the probability of an audit is strictly positive, collusion cannot be sustained if the firm and the inspector are either "too patient" or "too impatient". To understand why collusion is not sustainable when the players are extremely patient, note that as δ goes to one, their respective cumulated discounted payoffs from collusion are bounded above because of the positive probability of an audit.[15] That is because an audit drives both players' payoffs to zero forever. However, if the inspector and the firm never exchange a bribe, their cumulated discounted payoffs go to infinity go to infinity as δ goes to one. Hence, a collusive bribe does not exist for high values of δ. Clearly, this reasoning does not apply if $\lambda = 0$. Thus, in this case, as shown in Figure 2b, collusion is feasible if and only if the discount factor is sufficiently high, specifically, $\delta \geq \tilde{\delta} \equiv \frac{\mu_n r f - e_n}{\mu_n f} (< 1)$. The fact that for $\lambda = 0$, there is a non-empty interval of discount factors for which preemptive corruption is sustainable, implies that, by continuity of L and R, such a range of discount factors always exists for λ strictly positive as long as λ is sufficiently small. If λ is relatively high, collusion is unfeasible for every value of δ as shown in Figure 2c.

We next examine the range of values of w where preemptive corruption is feasible.

Proposition 3 *Suppose equation (3) holds with a strict inequality when $w = \mu_n f$. Then there exist finite values of the private benefit from engaging in the illegal activity $\underline{w} < \mu_n f$ and $\overline{w} > \mu_n f$ such that collusive preemptive corruption is feasible if and only if the firm's benefit is in $[\underline{w}, \overline{w}]$. Therefore, firms that obtain a large benefit from engaging in the illegal activity or those that obtain a small benefit from engaging in the illegal activity do not engage in preemptive corruption.*

Proof. See the Appendix ■

Figure 3 graphically illustrates the above result and shows how the feasibility of collusion depends on the firm's private benefit from the illegal activity assuming that p_t, p_g and $\lambda > 0$ are sufficiently low. The figure depicts how the difference between the left-hand

[15] Specifically, for a bribe B^c, the finite upper bounds on the payoffs are $\frac{B^c - \lambda p_t}{\lambda}$ for the inspector and $\frac{w - \lambda(1+p_g)f - B^c}{\lambda}$ for the firm.

side and the right-hand side of (4) (for $w < \mu_n f$) and (3) (for $w \geq \mu_n f$) change with w assuming that (3) holds when w is arbitrarily close to $\mu_n f$. To construct Figure 3, we rely on the fact that $\frac{1-(1-\lambda)\delta}{(1-\lambda)\delta} < G(\lambda, \delta)$, which implies the discontinuity at $w = \mu_n f$. Firms that will engage in preemptive bribery are those for which w is between $\underline{w}$ and $\overline{w}$ where $\underline{w} \equiv \lambda(1 + p_g)f + \lambda p_t + \frac{1-(1-\lambda)\delta}{(1-\lambda)\delta}(\mu_n r f - e_n)$. Clearly, $\underline{w} = 0$ is possible.

It is worthwhile to contrast this result with those from static models of bribery. In static models w has no effect on corruption because it is earned whenever the firm chooses $W = 1$, which is a necessary condition for corruption. Furthermore, its value does not influence the inspector's choice of monitoring intensity. The value of w matters in this game because the firm has to forego w when it is detected by the regulator after having bribed the inspector. Firms with higher values of the private benefit from engaging in the illegal activity have more to lose from detection and thus will refrain from bribing the inspector. This result is interesting in light of Svensson (2003) who studies a cross-section of Ugandan firms that pay bribes to public officials. He finds that officials often price discriminate by altering the size of the bribe according to the profitability of the firm. Similarly, Lubin (2003) finds that small firms engage in different forms of corruption as compared to large firms.

3.3 Comparative Statics

Our analysis of preemptive corruption yields comparative statics results that are summarized in Proposition 4 below. For the result in Proposition 4, we say that a change in a parameter makes collusive preemptive bribery more difficult if it reduces the set of collusive perfect equilibrium preemptive bribes. Therefore other things constant, a change in a parameter makes collusion more difficult if it leads to a smaller difference between the left-hand side and the right-hand side of (3) or (4).

Note that our approach to derive comparative statics results differs from that used in Besley and McLaren (1993), Mookherjee and Png (1995) or Samuel (2008). These authors focus on specific bribes, ex-post or preemptive, obtained from solving the Nash bargaining problem between the firm and the inspector. Instead we are interested in how the set of bribes that allow the two proponents to earn discounted expected payoffs in excess of their static game payoffs is affected by changes in important policy parameters. Our approach

is therefore similar to that employed by Buccirossi and Spagnolo (2006) in the context of leniency programs.

Proposition 4 *Suppose Assumption 2 holds and equation (3) is satisfied. Then, other things constant, the following holds.*

(i) If $\frac{\partial \mu_n}{\partial r}\frac{r}{\mu_n}$, the elasticity of monitoring intensity by the inspector (μ_n) to a change in his reward r, is sufficiently high, then an increase in r (the inspector's reward) facilitates preemptive corruption.

(ii) Increases in λ (the probability that corruption is detected), p_t or p_g (the fines for corruption) unambiguously make preemptive corruption more difficult.

(iii) An increase in λ and a decrease in either p_g or p_f that leaves expected fines, $\lambda[p_t + (1 + p_g)f]$ unchanged unambiguously makes preemptive corruption more difficult.

Furthermore, if Assumption 2 does not hold, but (4) is satisfied, then, other things constant, (ii) and (iii) hold. However, in this case, an increase in r unambiguously makes preemptive collusion more difficult.

Proof. See the Appendix. ■

From a policy standpoint, parts (ii) and (iii) in the above result suggest that, although the probability that corruption is detected and the penalties for corruption may be viewed as substitutes, increasing the probability of detection is more effective at deterring collusion than increasing the fines. In previous papers the probability of detecting corruption and the penalties from corruption are perfectly substitutable (i.e., an expected-fine-neutral change in these policy parameters has no effect on corruption). Thus, a policy maker can choose between increasing the penalties or the probability of detection. In our model, λ has a dynamic effect that is absent in static analyses because being detected implies foregoing a entire stream of positive payoffs.

The model's prediction regarding the effect of the inspector's reward on the feasibility of preemptive corruption may be sharpened slightly by writing the elasticity of intensity to r as a function of the parameters of the model. First, using the first and second order conditions

that characterize μ_n, we obtain $\frac{\partial \mu_n}{\partial r} = \frac{f}{e''(\mu_n)} > 0$. Furthermore, $rf = e'(\mu_n)$. Therefore, $\frac{\partial \mu_n}{\partial r}\frac{r}{\mu_n} = \frac{rf}{e''(\mu_n)\mu_n} = \frac{e'(\mu_n)}{e''(\mu_n)\mu_n}$. That is, the elasticity depend on the first derivative and second derivatives of the cost-of-effort intensity. There are simple specifications of the cost of effort function satisfying Assumption 1 for which the elasticity of optimal effort to r can be quite large - possibly sufficiently large for the result that an increase in r facilitates preemptive corruption to occur. For example, for every $a > 1$, the parametric form $e(\mu) = \frac{\mu^a}{1-\mu}$ satisfies all the conditions in Assumption 1. Furthermore, for every $a \in (1, 2)$, there exists a range of values of μ for which the ratio $\frac{e'(\mu)}{e''(\mu)\mu}$ is well above 1 (a necessary condition for an increase in r to facilitate corruption).

4 Welfare discussion

In our model, the feasibility of preemptive corruption has important welfare implications. To analyze the welfare effects of policy changes that affect the feasibility of corruption, consider an economy made up of firms that differ solely with respect to their w. For instance, suppose that there are a continuum of firms uniformly distributed on an interval $[0, w_{Max}]$, where $h > w_{Max} > \mu_n f$ holds. In any given period, the contribution to social welfare of a firm that chooses $W = 1$ and is monitored at intensity level μ is $w - h - e(\mu) < 0$, while a firm that chooses $W = 0$ has a contribution of zero.

Since there are multiple equilibria in the repeated game, in the discussion below, we assume that if preemptive collusive corruption is feasible given the firm's w, then the firm and the inspector collude by exchanging some bribe out of the feasible set of bribes. Otherwise, they play the SPNE of the one-shot game (see Proposition 1). When the assumptions in the statement of Proposition 3 are satisfied, collusion is feasible with some types of firms, but not all. Figure 4 shows the various behaviors that arise depending on the value of w. In the remainder of the paper, we refer to firms with $w < \mu_n f$ as *potential offenders* and those with $w \geq \mu_n f$ as *regular criminals* since they always choose the bad technology, i.e., $W = 1$.

The behavior of the two marginal types, $\underline{w}$ and $\overline{w}$, will be affected by changes in the penalty levels (p_g and p_t), the probability of detection λ or even r, the inspector's reward. In turn, such policy changes will impact social welfare. Based on Figure 4, overall social

welfare or total surplus is equal to

$$\begin{aligned} TS &= \int_{\underline{w}}^{w_{Max}} \left(\frac{w-h}{w_{Max}}\right) dw - \left(\frac{w_{Max}-\overline{w}}{w_{Max}}\right) e_n \\ &= \left(\frac{w_{Max}-\underline{w}}{w_{Max}}\right)\left(\frac{w_{Max}+\underline{w}}{2} - h\right) - \left(\frac{w_{Max}-\overline{w}}{w_{Max}}\right) e_n < 0. \end{aligned}$$

Straightforward calculations establish that the above expression is increasing in $\overline{w}$ and decreasing in $\underline{w}$. Intuitively any policy change that lowers the value of $\overline{w}$ makes corruption more difficult, while it does not affect the firm's choice of technology. That is because regular criminals choose $W = 1$ even when they do not bribe the inspector (Figure 4). However, lowering $\overline{w}$ increases monitoring costs because when preemptive corruption occurs, no monitoring effort is spent. On the other hand, a policy change that raises the value of $\underline{w}$ does have an impact on the firm's choice of technology. That is because making preemptive corruption unfeasible for potential offenders implies that they will choose $W = 0$ (Figure 4). Also, in this case, making corruption more difficult does not raise monitoring costs since the firm chooses the good technology when it does not collude with the inspector.

It is straightforward to show that an increase in either p_g, p_t, or both, will lower $\overline{w}$ and raise $\underline{w}$. Thus, in our model, strengthening anti-corruption policies by increasing, p_g, p_t, or both only has a positive welfare impact among potential offenders, that is, firms with a low w, while such policy changes have a negative welfare impact on regular criminals with a high w. The effect of an increase in r is different. Increasing r has an ambiguous effect on $\overline{w}$, while it unambiguously raises $\underline{w}$. Furthermore, raising the inspector's reward also causes an increase in e_n, the one-shot SPNE level of monitoring effort.

We therefore find that a change in penalties for corruption and a change in the inspector's reward for fining the firm impact welfare in different ways. On the one hand, contrary to an increase in p_g or p_t, raising the inspector's reward may increase the scope for socially beneficial corruption. On the other hand, it also increases monitoring efforts with firms at the top of the distribution of private benefits. Unfortunately, in our model, it is not possible to obtain a sharper prediction regarding the effectiveness of r at improving the welfare outcome without imposing additional structure on the effort cost function.

5 Ex-post bribery and preemptive corruption with public randomization

So far, we have focused on preemptive corruption and have ignored the possibility of ex-post corruption. Intuitively, it seems that ex-post corruption should never occur when preemptive corruption is feasible. Indeed, ex-post corruption requires costly effort while preemptive corruption does not require any effort. However, with ex-post corruption, in every period, the probability of being detected for corruption is lower than under preemptive corruption because a bribe is exchanged with probability μ only. This suggests that there may be ways in which the firm and the inspector can reduce the probability of getting caught by reducing the frequency of bribe exchange, while avoiding costly supervision effort. This would of course require that the firm and the inspector are able to observe the realization of a public signal.

Suppose such public randomization is feasible and that the firm and the inspector agree to the following. In every period, depending on the realization of the random public signal, the firm pays a bribe $B > 0$ or pays no bribe. Whether or not the firm pays a bribe, the inspector exerts no effort and does not report the firm. The signal is i.i.d. across periods and such that, with probability α, the firm pays a bribe and with probability $1 - \alpha$, it pays no bribe. Therefore, in every period, the probability that the firm and inspector are caught for bribery is equal to $\alpha\lambda$. Clearly, the inspector's expected payoff along the collusive path is given by

$$V_S^{c\alpha} = \frac{\alpha(B - \lambda p_t)}{1 - \delta(1 - \alpha\lambda)}$$

and the firm's expected payoff is given by

$$V_F^{c\alpha} = \frac{w - \alpha[B + \lambda(1 + p_g)f]}{1 - \delta(1 - \alpha\lambda)}.$$

By setting $B = b$ and $\alpha = \mu$, it is straightforward to see that the preemptive bribe scheme is equivalent to an ex-post scheme in terms of size of the expected bribe and probability of getting caught. However, the preemptive scheme involves zero effort and therefore it Pareto dominates the ex-post scheme. Furthermore, on a path with ex-post bribery, the inspector's incentive to deviate is no less than it is on a path on which the inspector and the

firm exchange a preemptive bribe conditional on the realization of the public randomization device.[16] Therefore, if ex-post corruption is sustainable, then so is preemptive corruption with a public randomization device.

Surprisingly, for a given $B > \lambda p_t$, it is possible to show that the inspector prefers $\alpha = 1$ to any $\alpha < 1$. Therefore, the inspector does not gain from the randomization scheme and, by transitivity, he would not gain either from engaging in ex-post bribery as opposed to simple preemptive bribery. Clearly, the firm prefers $\alpha = 0$, since this implies that no bribe is ever exchanged. More importantly, we show below that for each sustainable preemptive bribe, there exists a scheme with public randomization that is sustainable and Pareto dominates the simple preemptive corruption scheme with α implicitly set equal to one. The result is stated assuming $w \geq \mu_n f$. It is straightforward to show that a similar result holds when Assumption 2 is not satisfied.

Proposition 5 *Suppose Assumption 2 holds and assume $\lambda > 0$. Furthermore, suppose that equation (3) holds and let $\tau^c = \{(1, B^c, 0, 0)\}$ be a collusive stationary perfect equilibrium path. Then, there exists an α^c such that $0 < \alpha^c < 1$ and*

(i) the path $\tau^{c\alpha} = \{(1, \alpha^c, B^{c\alpha}, 0, 0)\}$, where $B^{c\alpha} = w - \lambda(1 + p_g)f$, can be supported as a perfect equilibrium,

(ii) $V_F(\tau^{c\alpha}) \equiv \frac{(1-\alpha^c)w}{1-(1-\alpha^c\lambda)\delta} = V_F(\tau^c)$,

(ii) $V_S(\tau^{c\alpha}) \equiv \frac{\alpha^c[w-\lambda(1+p_g)f-\lambda p_t]}{1-(1-\alpha^c\lambda)\delta} > V_S(\tau^c)$.

Therefore, preemptive corruption with public randomization yields a Pareto improvement over simple preemptive corruption.

Proof. See the Appendix ■

The proposition shows that whenever preemptive corruption arises without the use of public randomization, then there exists a collusive scheme with public randomization that is preferred by the inspector and no worse from the firm's standpoint. Using $\alpha < 1$ allows the corrupt pair to avoid the regulator's scrutiny with probability $1 - \alpha$. The collusive scheme

[16] A formal proof is available upon request from the authors. However, intuitively this occurs because the cost of effort is sunk under ex-post collusion. Thus, conditional on discovering evidence, the deviation payoff under ex-post bribery is greater than the deviation payoff under preemptive bribery with randomization. Thus, if the ex-post incentive constraint is satisfied, the preemptive collusive constraint under randomization is satisfied *a fortiori*.

described in the proposition is such that when a bribe is exchanged, it leaves no surplus to the firm. However, in periods in which bribery does not occur, the inspector earns a payoff of zero, while the firm earns w. The value of α is chosen so that the firm is indifferent between the scheme with public randomization and a simple preemptive bribe of B^c. For this value of α, the inspector earns a strictly higher payoff with public randomization than without it (and a bribe of B^c) and his incentive to deviate is no greater. Interestingly, there appears to be suggestive evidence of this type of randomized preemptive corruption in the case of the Gujurat trucking industry discussed earlier. Indeed, survey results by the Center for Electronic Governance (2002) indicate that inspectors at weigh station randomly select truck drivers to whom they charge a bribe.

6 Conclusion

Collusive contracts are, by definition, not legally enforceable and therefore they are easily susceptible to the hold-up problem. It has often been argued that this potential for hold-up creates disincentives for corruption and bribery. Our paper shows that whereas hold-up clearly prevents bribery between a firm and an inspector in a one-shot game, preemptive corruption may be supported as an equilibrium of an infinitely repeated game because the firm and the inspector can use the trigger strategy to punish deviations. Indeed, in our model, although the firm's hands are tied because it does not have a profitable deviation, repeated interactions can still enforce the collusive contract. Thus, in countries where bureaucrats have relatively strong control rights, preemptive bribery can still exist in the presence of hold-up.

In addition to providing an explanation for bribery in the presence of hold-up, our model also helps explain an empirical result that cannot be easily reconciled with the one-shot games of Mookherjee and Png (1995) and Samuel (2008). In Svensson's (2003) empirical analysis of bribery among firms in Uganda, the likelihood of paying a bribe as well as the size of the bribe differ across types of firms (i.e., highly profitable or less profitable firms). Similar findings are also discussed in Lubin's (2003) case study. In our model, firms that engage in preemptive bribery are not those that benefit the most, nor those that benefit the least from pursuing the illegal activity, but those in between. However, among those that

do engage in preemptive bribery, firms that stand to gain more from the illegal activity pay higher bribes. Thus, the incentives for bribery and the size of bribe depend on the magnitude of a firm's gain from the illegal activity as in Svensson's empirical study.

Our model offers insight into the anti-corruption policies suggested by Mookherjee and Png's (1995) model within the context of a repeated game. Similar to Mookherjee and Png (1995) and Samuel (2008), we examine the impact of the probability of detecting bribery, and the subsequent penalties on bribe givers and bribe takers, on the feasibility of corruption.[17] In their one-shot games the probability of detecting corruption and the penalties for corruption are perfectly substitutable. Therefore, an expected-fine-neutral change leaves the incentives for corruption unchanged. However, this is not true in the repeated game where an expected-fine-neutral change in these policy parameters unambiguously makes corruption more difficult. Consequently an increase in the probability of detection (λ) is more effective at deterring collusion than an increase in the penalties (holding expected fines constant). This result may be particularly useful in situations where political constraints make it difficult to increase penalties and where limiting repeated interactions between bureaucrats and firms is impossible because implementing a staff-rotation policy is too costly. Our model suggests that in such cases raising the probability of detection will be more effective in deterring corruption.

Although increasing penalties for corruption and the probability of detection make corruption more difficult, the welfare implications of such changes are ambiguous because of our result that the firms that are most likely to offer bribes are not those that benefit the most or the least from the illegal activity. In particular, when the penalties for bribery are raised, then monitoring is increased and bribery reduced among regular criminals. Since regular criminals always pursue the illegal activity irrespective of whether they pay a bribe, raising the penalty for bribery does not lower the harm associated with their illegal actions even though it generates higher monitoring costs. Thus, raising the penalties for bribery produces a net welfare loss from regular criminals. By contrast, bribery encourages some potential offenders to choose the illegal activity when they would otherwise not have done so. Thus, raising the penalties for bribery reduces both bribery and the illegal activity among

[17] This is in contrast with Besley and McLaren's (1993) model, in which no fines are imposed on bribe givers.

potential offenders, and produces a net welfare gain from these firms. Consequently, raising the penalties for bribery has an ambiguous effect on welfare.

Increasing the inspector's reward for truthful reporting has a similarly ambiguous effect on welfare. Indeed, on the one hand, increasing the inspector's reward discourages potential offenders both from engaging in corruption as well as from engaging in the illegal behavior, leading to a net welfare gain. On the other hand, it may also increase the fraction of regular criminals who offer a preemptive bribe thereby reducing inspection costs, which could lead to a net welfare gain. Thus, increasing the inspector's reward for truthful reporting may not lead to less corruption although it may increase welfare.

Finally, we show that if the firm and the inspector can condition their bribe exchange on the realization of a public randomization device, then should corruption occur, it will be preemptive rather than ex-post corruption. Moreover, alternating stochastically between periods of corruption and periods when no bribe is exchanged raises the payoff of both parties involved in bribery. These results have the following implications. First, they imply that evidence of a bribe may not exist every time the firm and the inspector interact, even though the pair engages in preemptive corruption. Second, our results suggest that if eliminating corruption is desirable, then the regulator may be justified in focusing on preemptive rather than ex-post corruption. Finally, they also suggest that when finding direct evidence of bribe exchange is costly, the regulator's resources may be better spent monitoring the inspector so as to insure inspections are carried out, thereby making preemptive corruption more difficult. This last implication is particularly relevant in the context of e-government programs that use computers to record inspection effort (Center for Electronic Governance 2002). Given the hold-up problem that characterizes one-shot interactions, our findings suggest that these policies could be effective in dealing with preemptive corruption insofar as they prevent inspector shirking. Indeed, forcing the inspector to spend monitoring effort increases his incentive to cheat on the preemptively collusive agreement. Although ex-post corruption may be feasible, we have shown that this form of collusion is more difficult to sustain, and thus easier to eradicate.

7 Appendix: Proofs

Proof of Proposition 1

We solve for the subgame perfect Nash equilibrium of the extensive form game in Figure 1 using backward induction. To this effect, note first that for every $B \geq 0$, if the inspector chose $\mu > 0$, successfully produced evidence, after which the firm offered $b > 0$, then the inspector will choose AR since this maximizes his expected payoff. Consequently, $b = 0$ is optimal for the firm. It follows that on the equilibrium path, the inspector will choose μ to maximize $\mu r f - e(\mu)$, that is, the inspector chooses $\mu = \mu_n$. But then, it is clear that $\mu = \mu_n$ is optimal for every $B \geq 0$. It thus follows that a firm that chose $W = 1$ will always choose $B = 0$. We have thus established that if the firm chooses $W = 1$, then it will offer $B = 0$. The inspector then chooses $\mu = \mu_n$ and produces evidence with probability μ_n. If the inspector finds evidence, then the firm chooses $b = 0$. Therefore, the firm's expected payoff from choosing $W = 1$ is equal to $w - \mu_n f$. Finally, the firm will choose $W = 1$ if and only if $w - \mu_n f \geq 0$.

We have thus shown that the paths described in the statement of Proposition 1 are the equilibrium paths depending on the value of w. Since in both cases backward induction has produced a unique solution, we conclude that the equilibrium is unique.

Proof of Proposition 2

The proof consists in determining a range of bribes B^c, which, if exchanged on a path $\tau^c = \{1, B^c, 0, 0\}$, yield collusive payoffs in the sense of Definition 1. Then, we provide conditions under which bribes in this range satisfy the necessary and sufficient incentive compatibility constraints (1) and (2).

Suppose that Assumption 2 holds. Consider the path $\tau^c = \{1, B^c, 0, 0\}$. We derive a lower bound $\underline{B}$ and an upper bound $\overline{B}$ for the bribe B^c so that τ^c is a collusive path in the sense of Definition 1. We have $V_S(\tau^c) = \frac{B^c - \lambda p_t}{1-(1-\lambda)\delta}$ and $V_F(\tau^c) = \frac{w - B^c - \lambda(1+p_g)f}{1-(1-\lambda)\delta}$. For τ^c to be collusive, Definition 1 requires

$$\frac{B^c - \lambda p_t}{1-(1-\lambda)\delta} \geq \frac{U_S^N}{1-\delta}$$

and

$$\frac{w - B^c - \lambda(1+p_g)f}{1-(1-\lambda)\delta} \geq \frac{U_F^N}{1-\delta}$$

with at least one strict inequality. Rearranging the above inequalities yields

$$B^c \geq \lambda p_t + [1-(1-\lambda)\delta]\frac{U_S^N}{1-\delta} \equiv \underline{B}$$

and

$$B^c \leq w - \lambda(1+p_g)f - [1-(1-\lambda)\delta]\frac{U_F^N}{1-\delta} \equiv \overline{B}$$

A path τ^c is collusive if and only if $\underline{B} < \overline{B}$ and $B^c \in [\underline{B}, \overline{B}]$. Straightforward calculations show that $\underline{B} < \overline{B}$ if and only if

$$w - \lambda(1 + p_g)f - \lambda p_t > \left(\frac{1 - (1 - \lambda)\delta}{1 - \delta}\right)(U_F^N + U_S^N). \tag{5}$$

If (5) holds, then a collusive stationary path exists. Assume (5) holds. Now, we wish to prove that if (3) holds, then there exists a path τ^c that satisfies both (1) and (2) and such that B^c, the bribe exchanged on τ^c is in the interval $[\underline{B}, \overline{B}]$. Then, we show that if (3) does not hold, then at least one of the following conditions is violated: (1), (2) or B^c is a collusive bribe.

Assume $B^c \in [\underline{B}, \overline{B}]$. To determine the subset of bribes B^c exchanged on a stationary path $\tau^c = \{1, B^c, 0, 0\}$ that also satisfy both (1) and (2), we first determine what constitutes an optimal deviation from τ^c for each of the inspector and the firm. First consider the firm. Suppose that in period t, it deviates from τ^c by choosing $W = 0$. In this case, it earns a discounted expected payoff of $0 + \frac{\delta}{1-\delta}U_F^N$. Suppose instead that it chooses $W = 1$ in period t, but deviates by offering $B^d \neq B^c$. In this case, the extensive trigger strategy specifies that the players follow the deviation by engaging in subgame perfect behavior for the continuation of the Inspection game at period t. Then, the firm and the inspector play the SPNE of the Inspection game in every such game from period $t + 1$ on. It is straightforward to show that this deviation earns the firm an expected discounted payoff of $U_F^N - B^d - \lambda(1+p_g)f + (1-\lambda)\frac{\delta}{1-\delta}U_F^N$ if $B^d > 0$ and $\frac{U_F^N}{1-\delta}$ if $B^d = 0$. Finally, it is clear that the firm cannot deviate by offering an ex-post bribe because ex-post bribes subgames are never reached on the path τ^c. It is then straightforward to show that under Assumption 2, for the firm, an optimal deviation consists in choosing $W = 1$ and offering $B^d = 0$. Therefore, under the extensive trigger strategy, the firm earns a discounted expected payoff equal to $\frac{U_F^N}{1-\delta}$ from an optimal deviation.

Now consider the inspector. Suppose the firm has offered τ^c in period t as prescribed by the initial path τ^c. If the inspector deviates by rejecting the bribe, then, under the extensive trigger strategy, it earns a payoff of U_S^N in period t and U_S^N in every period thereafter. Suppose that instead, it accepts the bribe but invests $\mu > 0$. It is clear that it is optimal for the inspector to invest μ_n in this case. Therefore, in period t, the inspector's payoff from such a deviation is equal to $B^c + U_S^N - \lambda p_t$ and it earns $(1 - \lambda)\frac{\delta U_S^N}{1-\delta}$ in the continuation of the game. Hence, the latter deviation is more profitable than the former if and only if

$$B^c + U_S^N - \lambda p_t + (1 - \lambda)\frac{\delta U_S^N}{1 - \delta} \geq \frac{U_S^N}{1 - \delta}, \tag{6}$$

which is equivalent to $B^c \geq \lambda p_t + \frac{\delta\lambda U_S^N}{1-\delta}$. Note that the right-hand side of this last inequality is strictly less than $\underline{B}$. Hence, since $B^c \geq \underline{B}$ by assumption, (6) holds.

To prove the *if* part of the statement of Proposition 2, suppose equation (3) holds and consider the path $\tau^* = \{1, B^*, 0, 0\}$ where $B^* = \lambda p_t + G(\lambda, \delta)U_S^N$. We now show that τ^* is a collusive perfect equilibrium path.

First we show that B^* is a collusive bribe, that is, (5) holds and $B^* \in [\underline{B}, \overline{B}]$. A straightforward calculation shows that

$$B^* \leq \overline{B} \iff w - \lambda(1+p_g)f - \lambda p_t \geq \left(\frac{1-(1-\lambda)\delta}{1-\delta}\right) U_F^N + G(\lambda,\delta)U_S^N,$$

which holds since (3) holds by assumption. We now show that $B^* > \underline{B}$. This last inequality follows from straightforward calculations.

$$\begin{aligned} B^* > \underline{B} \quad &\iff \quad G(\lambda,\delta) > \frac{1-(1-\lambda)\delta}{1-\delta} \\ &\iff \quad \frac{1-\delta+\lambda\delta^2-\lambda^2\delta^2}{\delta(1-\delta)(1-\lambda)} > 1-\delta+\lambda\delta \\ &\iff \quad 1-\delta+\lambda\delta^2(1-\lambda) > (1-\delta)\delta(1-\lambda)+\lambda\delta^2(1-\lambda) \\ &\iff \quad 1 > \delta(1-\lambda) \end{aligned}$$

which follows from $\delta, \lambda < 1$.

Now we show that τ^* is a perfect equilibrium path. If τ^* is a perfect equilibrium path, then both (1) and (2) must hold on τ^*. Since $B^* \geq \underline{B}$, $B^* > \lambda p_t + \frac{\delta\lambda U_S^N}{1-\delta}$. Thus, in any period t of τ^*, (1) can be written as

$$\frac{B^* - \lambda p_t}{1-(1-\lambda)\delta} \geq B^* - \lambda p_t + U_S^N + (1-\lambda)\frac{\delta U_S^N}{1-\delta}. \tag{7}$$

A straightforward calculation shows that the above equation is satisfied with equality by definition of B^*, so that (1) is satisfied on τ^*. Now, to show that (2) is satisfied on τ^*, we need to show that $V_F(\tau^*) = \frac{w-B^*-\lambda(1+p_g)f}{1-(1-\lambda)\delta} \geq \frac{U_F^N}{1-\delta}$ since an optimal deviation from τ^* yields a payoff of $\frac{U_F^N}{1-\delta}$ to the firm. However, this is equivalent to $B^* \leq \overline{B}$, which holds if and only if (3) holds.

Now consider the *only if* part of the statement. Suppose (3) does not hold. Then, from the above arguments, $B^* > \overline{B}$. However, it is straightforward to check that, by definition, B^* is the lowest bribe that the inspector will find profitable to accept without deviating by investing $\mu = \mu_n$. Indeed, no bribe B^c in the interval $[\underline{B}, B^*)$ is immune to a deviation whereby the inspector accepts B^c and then invest μ_n. Hence, $B^c \geq B^*$ is necessary on a stationary perfect equilibrium path. But then, from $B^c > B^* > \overline{B}$, we find that no such path satisfies the definition of a collusive path. Hence, we have shown that (3) is necessary and sufficient for a collusive stationary perfect equilibrium path to exist.

Now suppose that Assumption 2 does not hold. We wish to determine conditions under which there exists a collusive stationary path $\tau^c = \{1, B^c, 0, 0\}$ that is a stationary perfect equilibrium path. The proof of (ii) is similar to the proof of (i), therefore, we simply point out the relevant differences between the two proofs.

First, $U_F^N = 0$ since in the one-shot SPNE the firm chooses $W = 0$. Therefore, using reasoning similar to that employed in order to prove (i) above, the bribe B^c is a collusive

bribe if the following two conditions are satisfied (with a strict inequality in at least one of them):

$$\frac{B^c - \lambda p_t}{1-(1-\lambda)\delta} \geq 0 \iff B^c \geq \lambda p_t,$$

$$\frac{w - \lambda(1+p_g)f - B^c}{1-(1-\lambda)\delta} \geq 0 \iff B^c \leq w - \lambda(1+p_g)f.$$

Now we turn to incentive compatibility constraints. Under the extensive trigger strategy, choosing $W = 0$ is also the firm's optimal deviation from the initial path. This deviation earns the firm a payoff of zero. Hence, if B^c is collusive, the firm has no incentive to deviate. However, if the firm has conformed to the path and offered B^c, an optimal deviation by the inspector consists in choosing $\mu = \mu_n$. Then, under the extensive trigger strategy, no ex-post bribe is offered and thus the inspector earns $B^c + \mu_n rf - e_n - \lambda p_t$ from the deviation. Hence, the inspector's incentive compatibility constraint is given by:

$$\frac{B^c - \lambda p_t}{1-(1-\lambda)\delta} \geq B^c - \lambda p_t + \mu_n rf - e_n$$

which implies

$$B^c \geq \lambda p_t + \frac{1-(1-\lambda)\delta}{(1-\lambda)\delta}(\mu_n rf - e_n) \equiv \underline{B}' > \lambda p_t.$$

It follows that $\tau^c = \{1, B^c, 0, 0\}$ is a collusive stationary perfect equilibrium path if and only if

$$\underline{B}' \leq w - \lambda(1+p_g)f \iff w - \lambda(1+p_g)f - \lambda p_t \geq \frac{1-(1-\lambda)\delta}{(1-\lambda)\delta}(\mu_n rf - e_n),$$

where the second inequality is equation (4).

Proof of Proposition 3

To prove Proposition 3, we first show that when (3) is satisfied at $w = \mu_n f$, then there exists a range of values of w for which (4) is satisfied and a range of values of w for which (3) is satisfied. To this effect, substitute for $w = \mu_n f$ in (3) and notice first that $w = \mu_n f$ implies $U_F^N = 0$ and second, that the inspector's SPNE payoff in the Inspection game is $U_S^N = \mu_n rf - e_n$. Therefore, the only difference between (3) and (4) when $w = \mu_n f$ is the term in front of $\mu_n rf - e_n$. A straightforward calculation shows that that $\frac{1-(1-\lambda)\delta}{(1-\lambda)\delta} < G(\lambda, \delta)$, which implies that at $w = \mu_n f$, (3) implies (4). By continuity in w, both equations will hold in (distinct) neighborhoods of $\mu_n f$.

The remainder of the proof is straightforward. To the right of $w = \mu_n f$, (3) applies because Assumption 2 holds. In this case, recall that $U_F^N = w - \mu_n f$, while U_S^N does not depend on w. Letting L be the left-hand side of (3) and R its right-hand side, we show that $L - R$ decreases with w. We have $\frac{\partial L}{\partial w} - \frac{\partial R}{\partial w} = 1 - \frac{1-(1-\lambda)\delta}{1-\delta} < 0$ if and only if $\lambda\delta > 0$. Since (3) holds at $w = \mu_n f$ by assumption, it follows that (3) holds in an interval $[\mu_n f, \overline{w}]$, where $\overline{w}$ is the solution to $L - R = 0$. Finally, note that because $L - R$ is linear in w, $L - R = 0$

admits a unique and finite solution $\overline{w}$.

To the left of $w = \mu_n f$, (4) applies because Assumption 2 does not hold. Since in this case, the left-hand side of (4) is the only element that depends on w, it is straightforward to see that (4) will be satisfied in a half-open open interval of the form $[\underline{w}, \mu_n f)$, where $\underline{w} \equiv \lambda(1+p_g)f + \lambda p_t + \frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}(\mu_n r f - e_n) \geq 0$.

Proof of Proposition 4

Suppose Assumption 2 holds. Let L denote left-hand side of (3) and R its right-hand side. At a vector of the parameters such that (3) holds, a small change in a parameter x makes collusion more difficult if and only if $\frac{\partial(L-R)}{\partial x} < 0$.

Consider the effect of r. L does not depend on r. Hence, if R is increasing in r, then increasing r makes preemptive bribery more difficult. Differentiation of R with respect to r and straightforward calculations yield

$$\frac{\partial R}{\partial r} > 0 \iff \frac{\partial \mu_n}{\partial r}\frac{r}{\mu_n} < \left(\frac{1-\delta+\lambda\delta^2-(\lambda\delta)^2}{\delta(1-\delta)(1-\lambda)+\lambda\delta^2(1-\lambda)}\right) r, \tag{8}$$

where the term multiplied by r on the right-hand side of the second inequality is strictly greater than 1.

Next consider the effect of λ on $L-R$. A straightforward calculation yields

$$\frac{\partial(L-R)}{\partial\lambda} = -p_t - (1+p_g)f - \frac{\delta}{1-\delta}U_F^N - \frac{\delta^2(1-\lambda)^2+1-\delta}{\delta(1-\delta)(1-\lambda)^2}U_S^N,$$

which is clearly negative.

Finally, $\frac{\partial(L-R)}{\partial p_t} = -\lambda < 0$ and $\frac{\partial(L-R)}{\partial p_g} = -f < 0$. Hence, we have established (i) and (ii) in the statement of Proposition 4.

Now for (iii), note that simultaneous small changes $dp_t < 0$, $dp_g < 0$ and $d\lambda > 0$ that keep expected fine revenue constant must satisfy

$$\lambda dp_t + \lambda f dp_g + [p_t + (1+p_g)f]d\lambda = 0.$$

Hence, such a change will lead to $dL = 0$. However, $\frac{\partial R}{\partial\lambda} > 0$, while R depends neither on p_g, nor on p_t. Hence the proposed change results in $dR > 0$ and $d(L-R) = -dR < 0$, that is, it makes preemptive bribery more difficult.

When Assumption 2 does not hold, the comparative statics results follow directly from straightforward calculations and noting that the left-hand side of (4) (L) does not depend on r, while the right-hand side (R) increases with r. L is decreasing in p_t and p_g, but R does no depend on these parameters. Finally, L is decreasing in λ, while R is increasing in λ.

Proof of Proposition 5

To prove Proposition 5, we begin with (ii) and (iii). That is, we show that given a bribe B^c that is part of a collusive stationary perfect equilibrium path, there exists an α^c, $0 < \alpha^c < 1$ such that, setting $B^{c\alpha} = w - \lambda(1 + p_g)f$, on $\tau^{c\alpha}$, (ii) and (iii) are correct.

First note that the expected payoffs from conforming to the path with public randomization are $V_F(\tau^{c\alpha}) = \frac{(1-\alpha^c)w}{(1-(1-\alpha^c\lambda)\delta}$ and $V_F(\tau^c) = \frac{w-B^c-\lambda(1+p_g)f}{1-(1-\lambda)\delta}$. Then, define α^c as the unique solution in α to

$$\frac{(1-\alpha)w}{[1-(1-\alpha\lambda)\delta]} = \frac{w - B^c - \lambda(1+p_g)f}{1-(1-\lambda)\delta}.$$

The left-hand side of the equality is continuous and monotonically decreasing in α, so that α^c is uniquely defined. To show $0 < \alpha^c < 1$, note that $\alpha^c = 0$ implies that the left-hand side of the above equality is given by $\frac{w}{\delta}$. However, $\frac{w}{\delta} > \frac{w-B^c-\lambda(1+p_g)f}{1-(1-\lambda)\delta}$. This contradicts the definition of α^c. Similarly, $\alpha^c = 1$ implies that the left-hand side is equal to 0. Again, this contradicts the definition of α^c since $0 < \frac{w-B^c-\lambda(1+p_g)f}{1-(1-\lambda)\delta}$. Therefore (ii) holds by construction.

To prove (iii), we substitute for the value of α^c into $V_S(\tau^{c\alpha})$ and compare $V_S(\tau^{c\alpha})$ to $V_S(\tau^c)$. Straightforward calculations yield

$$V_S(\tau^{c\alpha}) > V_S(\tau^c) \iff w > -\frac{(1-\delta)[p_t + (1+p_g)f]}{\delta},$$

which always holds since $w \geq 0$. Therefore, we have shown (iii).

Now we turn to statement (i). Again, as in the case of preemptive corruption, under the extensive form trigger strategy, it is straightforward to show that the firm has no incentive to deviate from the path with public randomization. We therefore focus on the inspector's incentive compatibility constraints. To keep notation concise, on the path with the public randomization scheme, let the inspector's per-period payoff be $U_S^{c\alpha} = \alpha^c(B^{c\alpha} - \lambda p_t)$. There are two cases to consider. First, suppose that the firm offered the bribe $B^{c\alpha}$ based on the period-t signal. Then the inspector has two possible deviations. He may accept the bribe, but spend effort e_n and report the firm to the regulator with probability μ_n. In this case, his incentive compatibility constraint is given by

$$B^{c\alpha} - \lambda p_t + \frac{(1-\lambda)\delta U_S^{c\alpha}}{1-(1-\alpha^c\lambda)\delta} \geq B^{c\alpha} - \lambda p_t + \mu_n r f - e_n + \frac{(1-\lambda)\delta U_S^N}{1-\delta}. \tag{9}$$

The other possible deviation is for the inspector to turn down the bribe offer, but spend effort e_n and report the firm to the regulator with probability μ_n. With this deviation, there is no detectable corruption and the inspector's incentive compatibility constraint is given by

$$B^{c\alpha} - \lambda p_t + \frac{(1-\lambda)\delta U_S^{c\alpha}}{1-(1-\alpha^c\lambda)\delta} \geq \mu_n r f - e_n + \frac{\delta U_S^N}{1-\delta}. \tag{10}$$

The other case occurs when the firm does not offer a bribe, conforming to the period-t signal. Then, again, the inspector may spend effort e_n and report the firm to the regulator

with probability μ_n. In this case, his incentive compatibility constraint is given by

$$0 + \frac{\delta U_S^{c\alpha}}{1-(1-\alpha^c\lambda)\delta} \geq 0 + \mu_n rf - e_n + \frac{\delta U_S^N}{1-\delta}. \tag{11}$$

Since $\lambda < 1$, it is clear that (11) is a weaker condition than (9). Thus, if (9) holds, (11) is satisfied as well. Hence, the relevant constraints are (9) and (10). It is straightforward to show that (10) is always satisfied on a collusive path. Indeed, τ^c is collusive which implies that $\tau^{c\alpha}$ is also a collusive path. By definition of a collusive path $V_S^{c\alpha} = \frac{\alpha^c(B^{c\alpha}-\lambda p_t)}{1-(1-\alpha^c\lambda)\delta} \geq \frac{U_S^N}{1-\delta}$. The right-hand side of (10) is equal to $\frac{U_S^N}{1-\delta}$ and the left-hand side is strictly greater than $V_S^{c\alpha}$. Therefore, (10) holds.

Now consider (9). It is clear that, because (ii) and (iii) hold, (9) is implied by (7) with $B^* = B^c$, the inspector's incentive compatibility constraint on τ^c. Therefore, we have shown that if τ^c is a collusive stationary perfect equilibrium path, then so is $\tau^{c\alpha}$. This completes the proof of Proposition 5.

References

[1] Acemoglu, D. and T. Verdier (2000) "The choice between market failures and corruption," *American Economics Review*, 90, 194-211.

[2] Bardhan, P. (2005) "The economist's approach to the problem of corruption," *World Development*, 34, 341-348.

[3] Becker, G. and G. Stigler (1974) "Law enforcement, malfeasance and compensation of enforcers," *Journal of Legal Studies*, 3, 1-18.

[4] Besley, T. and J. McLaren (1993) "Taxes and bribery: The role of wage incentives," *Economic Journal*, 103, 119-141.

[5] Buccirossi, P. and G. Spagnolo (2006) "Leniency policies and illegal transactions," *Journal of Public Economics*, 90, 1281-1297.

[6] Center for Electronic Governance (2002) "Computerized interstate check posts of Gujarat State, India: A cost-benefit evaluation study," Indian Institute of Management, Ahmedabad.

[7] Choi, J-P. and M. Thum (2004) "The economics of repeated extortion," *Rand Journal of Economics*, 35, 203-223.

[8] Di Tella, R. and E. Schargrodsky (2003) "The role of wages and auditing during a crackdown on corruption in the city of Buenos Aires," *Journal of Law and Economics*, 46, 269-92.

[9] Jain, A. (2001) "Corruption: A review," *Journal of Economic Surveys*, 15, 71-121.

[10] Klochko, M. and P. Ordeshook (2003) "Corruption, cooperation and endogenous time discount rates," *Public Choice*, 115, 259-283.

[11] Lambert-Mogiliansky, A., M. Majumdar and R. Radner (2007) "Strategic analysis of petty corruption: Entrepreneurs ," *Journal of Development Economics*, 35, 203-223.

[12] Lubin, N. (2003) "Who's watching the watchdog?" *Journal of International Affairs*, 56, 43-56.

[13] Lambsdorff, J. Graf (2002) "Making corrupt deals: Contracting in the shadow of the law," *Journal of Economic Behavior and Organization*, 48, 221-241.

[14] Lambsdorff, J. Graf and M. Nell (2007) "Fighting corruption with asymmetric penalties and leniency," CEGE Discussion Paper No. 59, Georg-August-Universität Göttingen.

[15] Mailath, G. and L. Samuelson (2006) *Repeated games and reputation: Long-run relationships*. Oxford University Press.

[16] Mookherjee, D. and I.P.L. Png (1995) "Corruptible law enforcers: How should they be compensated?" *Economic Journal*, 105, 145-159.

[17] Motta, A. (2009) "Ex-ante and ex-post corruption," Marco Fanno Working Paper No. 94, University of Padua.

[18] Motta, M. and M. Polo (2003) "Leniency programs and cartel prosecution," *International Journal of Industrial Organization*, 21, 347-379.

[19] Polinsky, A.M. and S. Shavell (2001) "Corruption and optimal law enforcement," *Journal of Public Economics*, 81, 1-24.

[20] Samuel, A. (2008) "Preemptive collusion among corruptible law enforcers," *Journal of Economic Behavior and Organization*, Forthcoming.

[21] Strausz, R. (1997) "The delegation of monitoring in a principal-agent relationship," *Review of Economic Studies*, 64, 337-357.

[22] Svensson, J. (2003) "Who must pay bribes and how much? Evidence from a cross section of firms," *Quarterly Journal of Economics*, 118, 207-230.

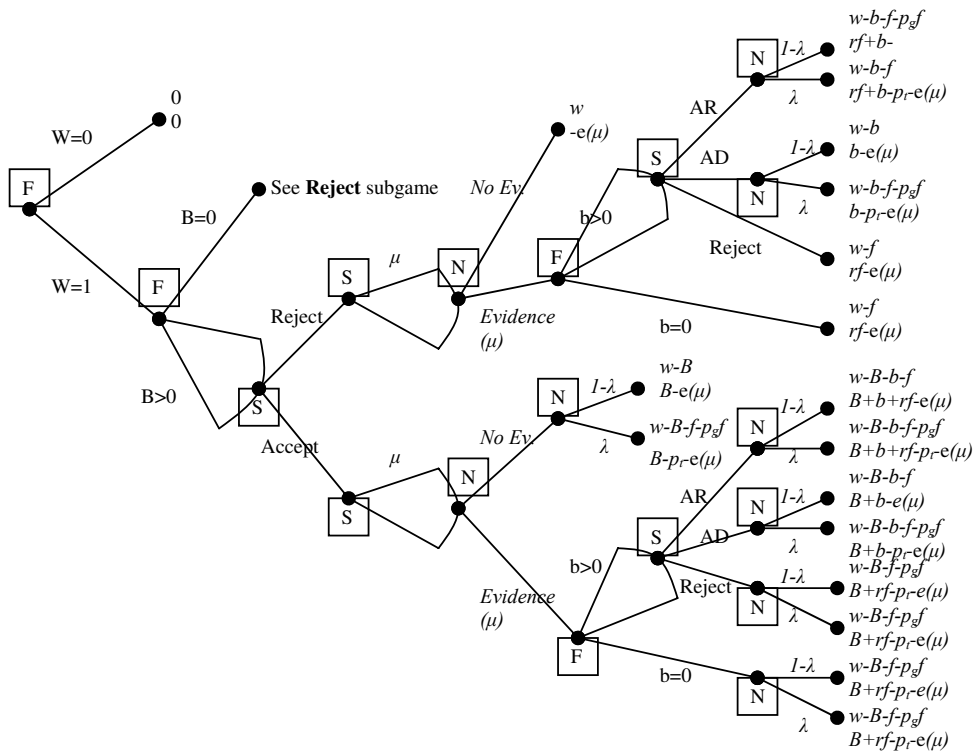

Figure 1: Extensive form of the Inspection game. N refers to a move by nature, F refers to a move by the firm and S refers to a move by the inspector.

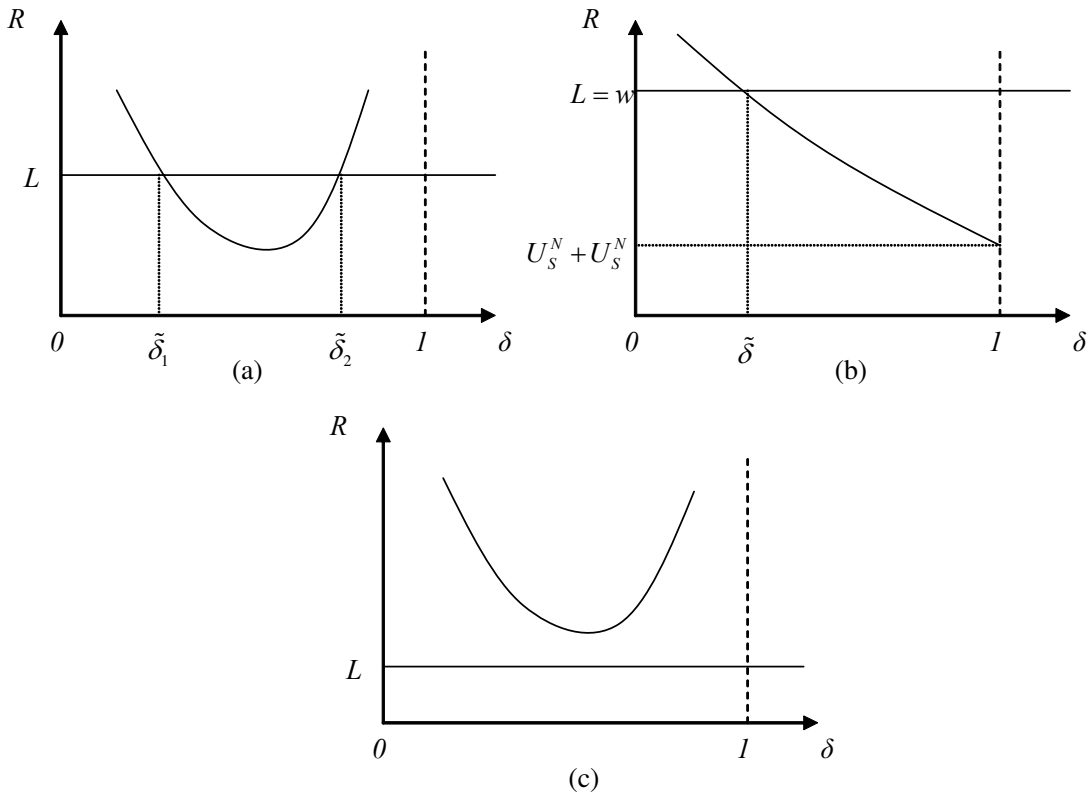

Figure 2: The range of δ's for which collusive preemptive collusion is feasible is the range for which the curve R (right-hand side of (3)) is below the line L (left-hand side of (3)) when Assumption 2 is satisfied. Each graph is drawn for a different value of λ.

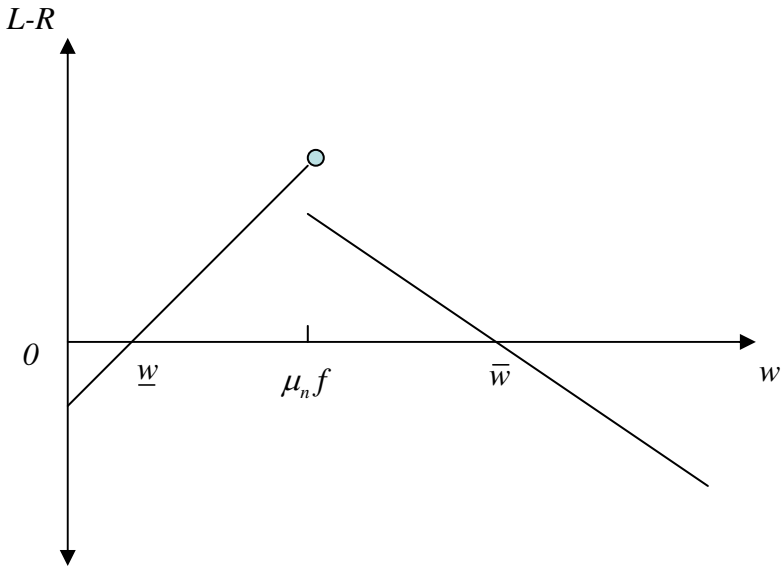

Figure 3: Illustration of Proposition 3. Range of w's for which either (3) or (4) holds assuming that (3) is satisfied when $w = \mu_n f$.

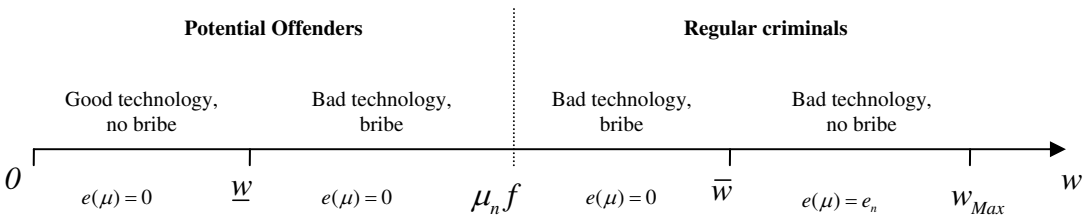

Figure 4: Typology of behavior and monitoring effort as a function of the firm's level of private benefit w.