Time to Come Clean? Disclosure and Inspection Policies for Green Production

Sang-Hyun Kim

Yale School of Management, Yale University, 135 Prospect Street, New Haven, CT 06511

sang.kim@yale.edu

February 2013

Abstract

We examine the interplay between two important policies that impact the environmental performance in a production setting: inspections performed by a regulator and noncompliance disclosure by a manufacturing firm. Expecting that a penalty will be levied once an inspection discovers noncompliance, the firm decides whether it should disclose a random occurrence of noncompliance. Anticipating this, the regulator determines the inspection frequency and the penalty amount that minimize environmental and social costs, performing either periodic inspections or random inspections. We study this problem by developing a novel analytical framework that combines the features from reliability theory and law enforcement economics. We find that, contrary to a common belief, a threat of increased penalty does not always lessen the need for costly inspections; there are situations where the regulator should invest in frequent inspections to complement the penalty. We also find that, counter to intuition, surprising the firm with random inspections is not always preferred to inspecting the firm periodically according to a set schedule.

Subject classification: Environment. Reliability: inspection. Games/group decisions: noncooperative.

Area of review: Environment, Energy, and Sustainability

Time to Come Clean? Disclosure and Inspection Policies for Green Production

Sang-Hyun Kim

Yale School of Management, Yale University, 135 Prospect Street, New Haven, CT 06511, sang.kim@yale.edu

February 2013

We examine the interplay between two important policies that impact the environmental performance in a production setting: inspections performed by a regulator and noncompliance disclosure by a manufacturing firm. Expecting that a penalty will be levied once an inspection discovers noncompliance, the firm decides whether it should disclose a random occurrence of noncompliance. Anticipating this, the regulator determines the inspection frequency and the penalty amount that minimize environmental and social costs, performing either periodic inspections or random inspections. We study this problem by developing a novel analytical framework that combines the features from reliability theory and law enforcement economics. We find that, contrary to a common belief, a threat of increased penalty does not always lessen the need for costly inspections; there are situations where the regulator should invest in frequent inspections to complement the penalty. We also find that, counter to intuition, surprising the firm with random inspections is not always preferred to inspecting the firm periodically according to a set schedule.

1 Introduction

Enforcement of environmental regulations is fraught with challenges. Not only do the regulatory bodies such as the Environmental Protection Agency operate under budget and resource constraints that limit their ability to reign on potential violators, they also face a nontrivial problem that renders enforcement particularly difficult: information acquisition. Unlike in many settings where a firm's private knowledge or action is reflected in performance outcomes that can be directly observed (for example, a user of equipment can estimate its reliability after experiencing downtimes during its operation), without direct consumers of environmental outputs, the indicators of a firm's regulation compliance are not readily available. Indeed, many environmental violations—especially those committed by small firms—go unreported because their individual output is not large and visible enough to trigger an alarm. Collectively, however, the sum of individual contributions may cause serious and potentially permanent damages to the environment (Kolstad 2011).

Under such a circumstance, the regulator who wishes to obtain information about a firm's compliance status would have no realistic choice other than to inspect the firm on site. Inspections are costly, especially because most violations occur sporadically in dispersed locations. These logistical difficulties are compounded by the fact that inspections are imperfect, in many cases due to the random nature of violations that occur despite the firms' best intentions. Such instances include: accidental release of untreated wastewater, excessive carbon emissions due to control system malfunctions, and toxic chemical spills following natural disasters. As Beavis and Walker (1983) put it, "Dischargers [of pollutants] are frequently unable to control with any great degree of accuracy the quantity and quality of wastes associated with specific levels of their productive activities." According to Malik (1993), "Pollution emissions by firms commonly depend on stochastic events such as equipment malfunctions, variations in input quality, and process upsets." In the presence of these uncertainties, an enforcement authority faces a challenge in devising cost-effective inspection strategies.

The tradeoff between cost of inspections and the social benefit of environmental preservation has been explored in the academic literature, starting from a broader context of law enforcements. One of the most well-known economic insights coming from this literature is that inspection intensity and the degree of sanction act as substitutes. That is, the regulator can save the cost of inspections without affecting a potential violator's behavior if the latter is threatened with a large penalty; the larger the penalty, the lower the intensity of inspections needed.[1] This intuitive notion of substitutability is regarded as fundamental, as evidenced by the following quote from an environmental economics textbook: "Increased monitoring activity by enforcement officials will have an effect similar to increased punishment levels... In theory, higher penalties can always substitute for lowered enforcement efforts" (Goodstein 2011, p. 286).

In this paper we reexamine this claim from a new perspective. In doing so, we develop a novel analytical framework that operationalizes the optimal enforcement decisions in the presence of stochastically evolving compliance states, inspired by the "inspection models" found in the theory of reliability. These model features are naturally built in a production setting where a manufacturing firm undergoes occasional environmental violations and restores compliance after each occurrence. We enrich this framework by adding an element of incentives: the firm's voluntary disclosure of noncompliance. That is, the firm makes a self-interested decision to either keep silence about an unintentional violation or disclose it to "come clean." Anticipating such a behavior, the regulator employs one of the two inspection policies: periodic inspections and random inspections. Under periodic inspections the regulator performs inspections according to a set schedule of constant intervals, whereas under random inspections the regulator randomizes the inspection intervals by sampling from a probability distribution. This new model framework allows us to address the following research questions that have been overlooked in the literature. Given that environmental violations occur randomly over time and that the firm may disclose noncompliance selectively, should the regulator perform periodic inspections or random inspections? How does the firm's opportunistic disclosure behavior impact the relationship between the two enforcement levers, inspection intensity and penalty?

We find that, contrary to the common belief found in the literature, inspection intensity and penalty are not necessarily substitutes; there are situations where the two act as complements. That is, the regulator may have to complement penalty with frequent inspections—stick and more stick instead of carrot and stick—in order to induce the desired behavior of the firm and minimize the social cost and environmental damage. We also find that, surprisingly, the regulator may find periodic inspections more cost-effective than random inspections. This is in spite of the fact that periodic inspections provide the firm with perfect knowledge about the inspection schedule whereas random inspections do not; even though better information may encourage the firm to act more opportunistically, such an adverse effect does not necessarily lead to a net loss in efficiency. We identify the conditions under which one inspection policy is preferred to the other, and offer managerial insights and policy recommendations.

2 Related Literature

Our model integrates the elements from two distinct streams of literature which, to the best of our knowledge, have never been put together in a single problem setting. They are: theory of reliability and economics of law enforcement. As we demonstrate, the ideas from each of these areas—developed in isolation over the years—bring new perspectives to the topic of environmental regulation once they are combined. Among the vast number of articles that have been published in these areas of research, in this section we survey the most relevant ones that inspired the model presented in this paper.

2.1 Theory of Reliability

Reliability theory concerns itself with evaluating performances of technological systems subject to random failures. See Barlow and Proschan (1996) and Rausand and Høyland (2004) for overviews of the foundations and applications of reliability theory. Among many models that have been proposed in the literature, the ones that have direct relevance to our problem are the models of inspection policies (Barlow and Proschan 1996, pp. 107-118). These models assume that the state of a system (i.e., whether the system is functioning or not) is normally invisible to a system operator. As a result, a system failure is not reported unless it is discovered by an inspection. An operator who wishes to detect a failure early faces a cost-benefit tradeoff because the chance of detection increases with the frequency of inspections but inspections are costly to perform. The models suggest optimal inspection schedules and strategies designed to balance this tradeoff, based on probabilistic representations of random failure processes.

A variant of the inspection models particularly relevant to our problem is the "intermittent faults" model, in which a system restores itself after each instance of repeated failures (Su et al. 1978, Nakagawa 2005, pp. 220-224). Thus the system alternates between "on" and "off" states stochastically over time without the intervention of an operator. Due to the random process, detection of a failure is not guaranteed. The operator's goal is to maximize the probability of detection by planning a frequent but cost-effective inspection schedule.

The inspection models—the model of intermittent faults in particular—provide the mathematical foundations that are well-suited for analyzing the problem of environmental regulation enforcements. First, a parallel can be made between random system failures and stochastic pollutions, since many environmental violations occur unexpectedly despite the firms' best intentions (e.g., when a pollution control system malfunctions). Second, in most cases inspections are needed to discover environmental violations, which are often hidden from the public's view. Third, the repeated nature of environmental violations and audits is captured in the inspection models, which evaluate long-term strategies for managing recurring failures. What lacks in the inspection models is the dimension of incentives, which we discuss next.

2.2 Economics of Law Enforcement

The basic premise of the inspection models, namely that constant monitoring is prohibitive because inspections are costly, was recognized in the economics literature by Becker (1968). In his model of probabilistic law enforcement, an enforcement authority combines random audits with sanctions (e.g., fines or imprisonment) in order to maximize the probability of apprehending the violators. A sanction influences the potential violator's behavior, an element missing in the inspection models. Becker's seminal work has been extended in many directions, notably by Kaplow and Shavell (1994) who incorporate voluntary disclosure. They consider individuals who have committed a socially harmful act but have an option to self-report it to the authorities. This feature is especially relevant to our model, as we incorporate self-disclosure of noncompliance .

The ideas introduced in these papers have been reinterpreted in the context of environmental regulations; see Cohen (1999) for a survey. Of particular relevance to this paper are the articles that consider enforcements in the presence of stochastic emissions, including Beavis and Walker (1982), Malik (1993), and Innes (1999). The messages are largely consistent with those from Kaplow and Shavell (1994) and others in the law enforcement literature, although idiosyncrasies exist. Russell (1990) is one of the few papers that feature a Markov process in an environmental audit context, but the focus is on random errors in audits instead of random occurrences of noncompliance. In a recent article, Toffel and Short (2011) establish hypotheses based on the theoretical predictions from these works as they empirically test whether the firms that self-report violations also self-police their operations to reduce emissions.

Despite some commonalities shared with these works, this paper distinguishes itself in a number of important aspects. First, we focus on environmental violations that occur unintentionally, driven by a random process. These types of violations are common in practice but they have received relatively little attention in the literature. Second, the way we model uncertainty is more realistic than those of Malik (1993), Innes (1999), and others. Unlike the stylized representations of random emissions found in these works, our model features a Markov chain of compliance/noncompliance states that alternate over time. Third, in contrast to Kaplow and Shavell (1994) and others who represent a random audit simply as a single probability measure, we bring precision to the mechanics of audits by modeling repeated inspections that may or may not detect noncompliance. Finally, the firm in our model decides not only whether noncompliance should be disclosed but also *when* it should be disclosed. This timing aspect has been ignored in the literature, and our model brings new perspectives by explicitly taking it into account.

2.3 Sustainable Operations

In recent years various issues of sustainability have drawn interests among the researchers in operations management (OM). Some have investigated traditional topics in environmental economics from operational perspectives (Sunar and Plambeck 2011, Drake et al. 2012, Alizamir et al. 2012). Others have studied new topics, including: adoption of green technology (Lobel and Perakis 2011, Avci et al. 2012); generating value over the product life-cycle (Lee 2011, Agrawal et al. 2012); configuring supply chains to reduce carbon footprints and be socially responsible (Caro et al. 2012, Cachon 2012, Guo et al. 2012). Although not framed as sustainability issues, Kim et al. (2010) and Kim and Tomlin (2012) share similarities with this paper in the modeling approaches based on reliability theory and the common theme of managing low-probability, high-consequence events.

More closely related to this paper are a number of recent articles that discuss audits and information disclosure in environmental regulations. Kalkanci et al. (2012) study the environmental impacts of voluntary and mandatory disclosure rules applied to supply chains. Jira and Toffel (2012) identify the conditions under which supply chain parties share information about greenhouse gas emissions. Plambeck and Taylor (2010) is one of the few articles in the OM literature that import ideas from the law enforcement economics literature, but they focus on issues that exist in a competitive market. Plambeck and Taylor (2012) study the dynamics that arise when firms make efforts to evade auditing, the topic we do not address in this paper. However, our findings on voluntary disclosure complement those in Plambeck and Taylor (2012). In sum, this paper is uniquely positioned in terms of research focus, insights, and the model features that bridge the gap between distinct areas of research.

3 Model

3.1 Overview

A firm ("he") produces a good and sells it to consumers. Production requires use of an environmentally harmful substance ("pollutant"), which may be emitted to the environment. A regulator ("she") is responsible for enforcing the environmental regulation. Both parties are risk neutral. Production and sales start at time zero and last over an infinite horizon. Production runs continuously unless it is temporarily suspended (more on this later). For the duration in which production is suspended, the firm loses a sales opportunity.

At any given moment the firm is in one of two states: the firm is said to be *in compliance* if pollutant emission is blocked, whereas it is *in noncompliance* if the pollutant is being emitted. These two states last for random amounts of time, and they alternate stochastically over time at constant transition rates.[2] The firm has complete visibility to the state but the regulator does not. There are two ways in which noncompliance is *reported* to the regulator: either the regulator discovers noncompliance after she performs an inspection (*detection*) or the firm preemptively informs it (*disclosure*). Note that we do not consider the firm's willful violations to save costs or attempts to evade inspections, the subjects beyond the scope of this paper. We assume that an inspection reveals the state immediately and that noncompliance is costlessly verified once it is disclosed.[3] Once noncompliance is reported, the regulator learns its start and end times (which together define the boundaries of the reported noncompliance *episode*) as she receives details of a violation while being kept informed of the progress made in restoring compliance. If the firm does not disclose a noncompliance episode, it may or may not be detected by the regulator because the episode lasts for a random amount of time; detection occurs if an inspection arrives before the episode concludes, whereas detection does not occur if the episode falls completely within the interval between two successive inspections. See Figure 1 for illustrations.

Production and inspections are suspended immediately after noncompliance is reported and remain suspended until compliance is restored. (If this assumption is violated, suboptimal outcomes arise in which unnecessary inspections are performed and avoidable emissions accumulate.) Both activities resume upon compliance restoration. Because a noncompliance episode may be reported after some time has passed since it started, as is the case when it is reported via detection, an episode may consist of two successive portions: *unsuspended noncompliance* and *suspended noncompliance*. Pollutant is emitted in the first portion but not in the second portion, during which there is no production output.

The regulator utilizes two enforcement levers: inspections and penalty. The latter is imposed upon discovery of noncompliance. Before time zero the regulator sets the inspection frequency and

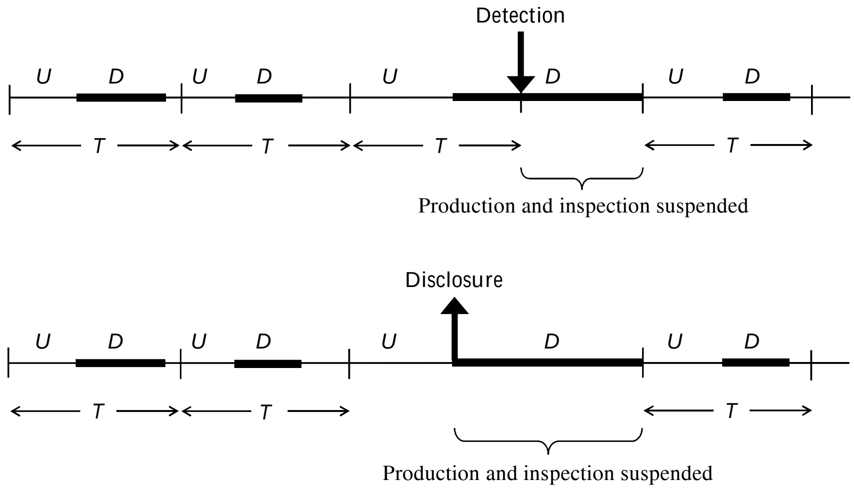

Figure 1: Illustrations of detection and disclosure instances. Compliance and noncompliance states alternate over time, while the regulator performs an inspection every T time units unless it is suspended due to a noncompliance report via detection or disclosure. The inspection process restarts upon compliance restoration. The independent and identically distributed durations of compliance and noncompliance episodes are denoted by U and D, respectively.

the penalty amount, subsequently announcing them to the firm. The regulator does not update these decisions once they are announced. In response, the firm devises a decision rule which specifies how much information should be disclosed preemptively. Thus, we seek a subgame perfect equilibrium of this two-stage game in which the regulator moves first as the Stackelberg leader. The inspection and disclosure policies employed by the regulator and the firm are detailed in §3.4. The regulator makes decisions to maximize the long-run average social welfare, whereas the firm makes decisions to maximize his long-run average profit.

3.2 Compliance State Transitions

If the firm is in compliance during production or if production is suspended following a report of noncompliance, no pollutant is emitted to the environment. (While zero emission is a simplifying assumption, relaxing it does not significantly impact the results.) On the other hand, if the firm is in noncompliance but production continues ("noncompliant production"), the pollutant is emitted at a constant rate.

The compliance and noncompliance states alternate as a two-state continuous-time Markov chain with the rates λ and μ, respectively. Hence, the firm stays in the compliance state for an exponentially distributed amount of time with mean $1/\lambda$ and in the noncompliance state for an exponentially distributed amount of time with mean $1/\mu$. The rate μ is interpreted as the capacity to restore compliance which the firm has installed before production starts. We assume that $\lambda \ll \mu$, i.e., noncompliance is a rare occurrence. Initially at time zero, the firm is in compliance. The probability that the firm starts in the compliance state at time zero and ends in the noncompliance state at time t is (Nakagawa 2005, p. 221)

$$\theta(t) \equiv \frac{\lambda}{\lambda+\mu}(1-e^{-(\lambda+\mu)t}). \tag{1}$$

We assume that the transition rates λ and μ are exogenously given and unaffected by managerial interventions.

3.3 Production and Economics

In order to highlight the dynamics arising from stochastic compliance state transitions and the inspection/disclosure policies that depend on them, in our model we keep the representation of the production process to a minimum. It is assumed that the demands for the firm's good are deterministic and that they arrive at a constant rate, normalized to one. Each unit of demand triggers a production order (i.e., make-to-order production). Hence, production runs at a constant rate equal to one unless it is suspended due to a report of noncompliance. Production and delivery lead times are negligible and there is no fixed setup cost. We assume that all unmet demands due to suspended production are lost and that no inventory is held by the firm.

The firm earns revenue r for each unit sold. To focus on the relationship between the firm and the regulator, we do not explicitly model the consumer purchasing behavior and assume that they make zero surplus for the units they acquire. Production cost is normalized to zero. If the firm continues production while not being in compliance, pollutant is emitted at a constant rate and causes environmental damage valued at h per unit time. Damage is avoided if production is suspended following a report of noncompliance. The value h reflects both immediate and long-term impacts on the environment, and we assume that it is higher than the opportunity cost of lost sales, i.e., $h > r$. This assumption implies that, from the society's standpoint, pollution prevention takes precedence over revenue generation.

The regulator incurs a fixed cost χ each time she performs an inspection. We assume that the cost of performing an inspection is sufficiently small so that the condition $\chi < (h-r)\lambda/\mu^2$ is satisfied.[4] This condition ensures that the regulator has an incentive to perform inspections; if it is violated, the firm may find inspections too costly to justify the potential social benefit. The firm is liable to the penalty κ, which is levied by the regulator if it is discovered that the damage to the environment has already been done. Therefore, the firm pays the penalty if the regulator's inspection detects noncompliance or if the firm is found to have been in noncompliance at the time of his disclosure. By contrast, the firm is exempt from the penalty if he discloses noncompliance as soon as it occurs, because the immediate suspension of production followed by such an action leaves no environmental damage.

We assume that there is an upper bound on κ whose value is exogenously given. This rules out an unrealistic scenario where the regulator imposes an arbitrarily large penalty on the firm. A similar assumption is commonly found in the literature (e.g., Kaplow and Shavell 1994, Innes 1999) and is often justified on the grounds that financial and legal restrictions such as the bankruptcy risk limit the penalty size. For notational convenience we specify the upper bound as $r\delta$, where δ is the normalized maximum penalty.[5] Henceforth the condition $\kappa \leq r\delta$ is referred to as *maximum penalty constraint*. In addition, we make the following mild technical assumption to further rule out unlikely situations: $\delta < \widehat{y}/\mu$ where $\widehat{y} > 0$ is the unique solution to the equation $\frac{y/2}{y/\ln(1+y)-1} = \frac{\mu}{\lambda} - 1$. When combined with the earlier assumption $\lambda \ll \mu$, this condition is easily satisfied by most reasonable values of δ; in practical scenarios the condition is effectively equivalent to $\delta < \infty$.[6]

3.4 Inspection and Disclosure Policies

We assume that the regulator employs one of two inspection policies: *periodic inspections* and *random inspections*. In both cases the regulator performs inspections every T time units—a variable we call *inspection interval*—unless an inspection is suspended due to a noncompliance report, in which case the inspection process restarts upon compliance restoration. Under the random inspection policy T is a random variable with mean τ, i.e., $E[T] = \tau$. Under the periodic inspection policy, on the other hand, $T = \tau$ since T is deterministic. The regulator sets τ in both cases, thus determining the *inspection frequency* $\nu \equiv 1/\tau$. (Note that, despite the constant value $T = \tau$ under the periodic inspection policy, in general the realized intervals are not uniformly spaced because an inspection may be suspended for a random amount of time following a report of noncompliance. See Figure 1 for illustrations.) We assume that T under the random inspection policy is exponentially distributed. This assumption is made in order to maximize the contrast between random and periodic inspections. Because of memorylessness, the firm facing random inspections with exponential intervals has identical outlook of the future at each point in time; the "surprise factor" of random inspections is maximal under this policy, since the timing of the past inspection is irrelevant in predicting when the next inspection will take place. This is in sharp contrast to the case of periodic inspections, since under this policy the past inspection informs the firm with perfect knowledge about the timing of the next inspection; in this case, the surprise factor is zero.

The firm is said to practice *nondisclosure* if he never discloses noncompliance. Similarly, *full disclosure* means that the firm discloses all occurrences of noncompliance, while *partial disclosure* means only selected ones are disclosed. The values of enforcement and economic parameters determine which practice is adopted by the firm. The simplest choice is between full disclosure and nondisclosure:

either disclose all or nothing. Such a binary decision makes sense in some situations, but in general admitting partial disclosure improves the firm's performance.

If the firm were to adopt partial disclosure, an important factor in deciding whether to disclose a particular occurrence of noncompliance is *when* it occurs. This is because there is a high chance that the firm will escape detection if the next inspection is to arrive far into the future. In such a case he is willing to take a chance and not disclose noncompliance, since it is likely that compliance will be restored before the next inspection. By contrast, such willingness is lower if noncompliance occurs shortly before the next inspection; in this case the firm may be better off disclosing it preemptively to avoid being detected and penalized. This reasoning suggests that the natural disclosure policy in the presence of such a "horizon effect" is of the threshold kind. We formally state the two disclosure policies described so far.

Definition 1 *Under the* binary disclosure policy*, the firm either discloses every occurrence of noncompliance at its onset or does not disclose any. Under the* threshold disclosure policy*, the firm discloses an occurrence of noncompliance at its onset if and only if the time remaining until the expected arrival of the next inspection is smaller than or equal to* $s \in [0, \tau]$.

We call the time interval of length s in the definition *disclosure window*. By definition this window never exceeds the expected inspection interval; hence it satisfies $0 \leq s \leq \tau$. Under the threshold disclosure policy the firm discloses only the noncompliance episodes that start within the disclosure window, situated in the later portion of the inspection interval. Hence, the larger the window, the more noncompliance episodes disclosed. Note that the threshold disclosure policy includes nondisclosure and full disclosure as special cases, each corresponding to $s = 0$ and $s = \tau$, respectively. We assume that the firm employs either the binary or threshold disclosure policy, choosing between the two depending on whether he is subject to periodic or random inspections. (As we show later, the threshold disclosure policy does not always present an advantage over the binary disclosure policy.)

An important part of Definition 1 is that the firm is assumed to disclose noncompliance immediately after it occurs. This is intuitive because the firm may avoid the penalty with certainty by not postponing disclosure. To streamline exposition we prescribe the immediate disclosure as part of the definitions of disclosure policies, noting that it is in fact optimal under both random and periodic inspections (proof is found in the accompanying Technical Appendix).

3.5 Decisions and Objectives

Before time zero, given the announced penalty $\kappa \geq 0$ and the inspection frequency $\nu > 0$ (or equivalently the mean inspection interval $\tau > 0$), the firm sets the disclosure policy by choosing between the threshold policy and the binary policy. Under the threshold policy, the firm sets the disclosure window size s. Under the binary policy, the firm chooses either nondisclosure or full disclosure. The firm's objective is to maximize his long-run average profit. Anticipating this response, the regulator chooses the values of κ and τ that maximize the long-run average social welfare subject to the maximum penalty constraint $\kappa \leq r\delta$.[7]

Let $\overline{I}$, $\overline{R}$, and $\overline{B}$ denote the long-run averages of the following performance measures, in the presented order: the number of inspections performed, the cumulative duration of suspended noncompliance, and the cumulative duration of unsuspended noncompliance. Under the assumptions outlined above, the long-run average social welfare is then equal to $r(1-\overline{R}) - \chi\overline{I} - h\overline{B}$: the firm earns the revenue r per unit time unless production is suspended, which lasts $\overline{R}$ per unit time in the long-run; the regulator incurs the fixed cost χ each time she performs an inspection, doing so $\overline{I}$ times per unit time in the long-run; the damage valued at h is done to the environment while noncompliant production lasts $\overline{B}$ per unit time in the long-run. Note that the penalty κ does not appear in the social welfare function because it is the amount of a transfer between the firm and the regulator that cancels out within the social boundaries.

From the expression above it is clear that maximizing the long-run average social welfare is equivalent to minimizing the long-run average social cost $\overline{C} \equiv \chi\overline{I} + r\overline{R} + h\overline{B}$, the convention we adopt in the remainder of the paper. Similarly, the firm's profit-maximization problem is formulated as the equivalent cost-minimization problem, $\overline{\Psi}$ denoting his long-run average cost. The expressions for $\overline{C}$ and $\overline{\Psi}$ are evaluated in the next section, where we characterize the equilibria under the periodic and random inspection policies.

4 Equilibria Under Periodic and Random Inspections

In this section we characterize the equilibria that emerge under the periodic and random inspection policies. In each case we evaluate the long-run average performance measures $\overline{I}$, $\overline{R}$, $\overline{B}$, and $\overline{\Psi}$ that define the objective functions, derive the firm's optimal response to the regulator's announcement of the penalty amount and the inspection frequency, and solve the regulator's social cost minimization problem. The comparisons of equilibria are presented in §5.

4.1 Periodic Inspections

Recall that the inspection interval T under periodic inspections is deterministic and equal to the constant τ. Before characterizing the equilibrium, we first develop the expressions for performance measures when the firm employs the threshold disclosure policy in response to periodic inspections.

4.1.1 Performance Measures Under Threshold Disclosure Policy

The most convenient unit of analysis in evaluating the performance measures under periodic inspections is an *inspection cycle*. An inspection cycle is the time interval that contains at most one inspection performed, beginning in the compliance state and concluding with three possible endings that leave the firm back in the compliance state: (i) the cycle ends τ time units after the start when the regulator arrives for an inspection and finds the firm in compliance; (ii) the cycle ends when compliance is restored following the regulator's detection of noncompliance at τ time units since the start; (iii) the cycle ends when compliance is restored following the firm's disclosure of noncompliance. Note that these are the only possible endings of an inspection cycle. Because of memorylessness of the exponential compliance and noncompliance durations, the start and end times of an inspection cycle mark regeneration epochs. Therefore, an inspection cycle forms a renewal (Heyman and Sobel 1982, p. 184; Tijms 2003, p. 40). Note also that an inspection cycle may end at, before, or after τ time units since the start, because it takes a random amount of time for compliance to be restored after detection or disclosure. Thus the length of an inspection cycle—denoted by X—is random, and it is to be distinguished from the constant inspection interval $T = \tau$.

Depending on whether the firm is in compliance at the beginning of the disclosure window (at time $\tau - s$ since the cycle start), under the threshold disclosure policy the three outcomes above are further divided into five cases, referred to as Case 1a, Case 1b, etc. See Figure 2 that illustrates these cases.

1. The firm is in compliance at time $\tau - s$ since the cycle start. Moreover: (a) compliance lasts until or after time τ, at which the current inspection cycle ends with an inspection (but no detection); (b) noncompliance starts before time τ and at that moment the firm discloses the state, and subsequently the current inspection cycle ends when compliance is restored.

2. The firm is in noncompliance at time $\tau - s$ since the cycle start. Moreover: (a) noncompliance lasts until or after time τ, at which the regulator's inspection detects noncompliance and subsequently the current inspection cycle ends when compliance is restored; (b) compliance is restored before time τ and it lasts until or after time τ, at which the current inspection cycle ends with an inspection (but no detection); (c) compliance is restored before time τ but it is followed by a transition to noncompliance before time τ which is disclosed by the firm, and subsequently the current inspection cycle ends when compliance is restored again.

Note that in all cases at most one new noncompliance episode occurs within the disclosure window because, under the threshold disclosure policy, the firm always discloses the first of such occurrences (see Cases 1b and 2c in Figure 2) and subsequently the current inspection cycle ends as soon as compliance is restored. As a result, the five cases described above form a complete list of categories

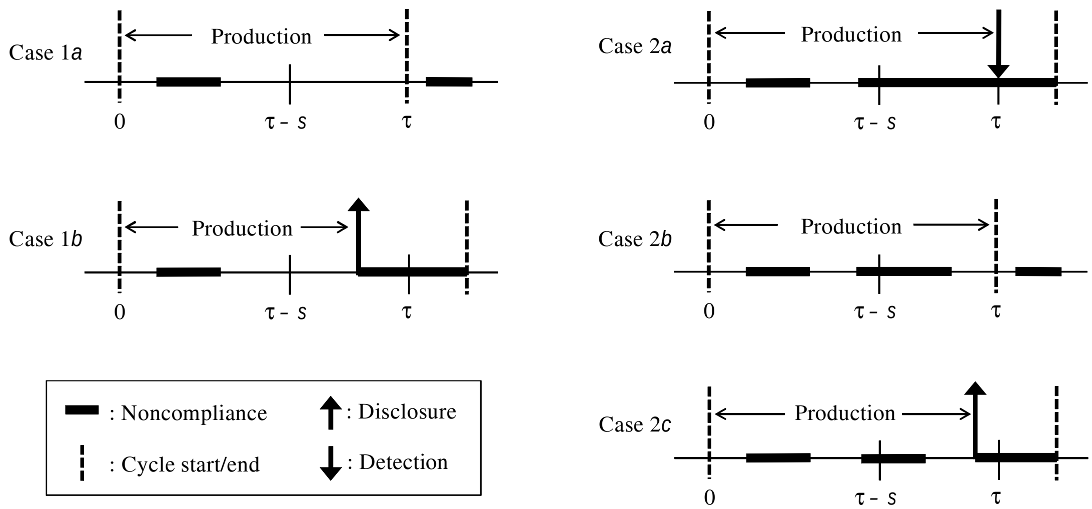

Figure 2: Five possible inspection/detection/disclosure outcomes that may arise in an inspection cycle under the threshold disclosure policy combined with periodic inspections. Noncompliance episodes are denoted by thick horizontal lines. Note that any realization of state transitions is possible before time $\tau - s$, provided that the state starts with compliance at time 0 and ends with either compliance (Cases 1a and 1b) or noncompliance (Cases 2a, 2b, and 2c) at time $\tau - s$. (For brevity, only one noncompliance episode appears in the figures.)

for all possible inspection/detection/disclosure outcomes that may arise in an inspection cycle. In addition, as illustrated in Figure 2, production lasts either until the firm discloses noncompliance or for exactly τ time units in case no disclosure is made.

With the complete list of stochastic outcome categories specified, we are now in a position to evaluate the long-run average performance measures. As an intermediate step, we first compute the probabilities of each case listed in Figure 2 and then evaluate the following quantities: (i) $E[X]$, the expected length of an inspection cycle; (ii) $E[I]$, the expected number of inspections performed in an inspection cycle; (iii) $E[R]$, the expected duration of suspended noncompliance in an inspection cycle; (iv) $E[B]$, the expected cumulative duration of unsuspended noncompliance in an inspection cycle; (v) $E[\Psi]$, the firm's expected cost in an inspection cycle. The results are summarized as follows.

Lemma 1 *(i)* $E[X] = \tau - s + \left(\frac{1}{\lambda} + \frac{1}{\mu}\right)(1 - e^{-\lambda s}) + \left(\frac{1}{\mu} - \frac{\lambda+\mu}{\mu(\mu-\lambda)}(e^{-\lambda s} - e^{-\mu s})\right)\theta(\tau - s)$*; (ii)* $E[I] = e^{-\lambda s} + \frac{\lambda}{\mu-\lambda}(e^{-\lambda s} - e^{-\mu s})\theta(\tau - s)$*; (iii)* $E[R] = \frac{1-e^{-\lambda s}}{\mu} + \frac{\mu e^{-\mu s} - \lambda e^{-\lambda s}}{\mu(\mu-\lambda)}\theta(\tau - s)$*; (iv)* $E[B] = \frac{\lambda}{\lambda+\mu}(\tau - s) + \left(\frac{1-e^{-\mu s}}{\mu} - \frac{1}{\lambda+\mu}\right)\theta(\tau - s)$*; (v)* $E[\Psi] = r\frac{1-e^{-\lambda s}}{\mu} + \left(r\frac{\mu e^{-\mu s} - \lambda e^{-\lambda s}}{\mu(\mu-\lambda)} + \kappa e^{-\mu s}\right)\theta(\tau - s)$.

All proofs are found in the Appendix. Note that the function $\theta(t)$ appearing in the lemma is given by (1). Since each inspection cycle forms a renewal, we apply the renewal-reward theorem (Tijms 2003, p. 41) to compute the long-run averages by forming a ratio between each performance measure in (ii)-(v) of the lemma and $E[X]$ in (i). For instance, the long-run average number of inspections performed is evaluated as $\overline{I} = E[I]/E[X]$. The remaining measures $\overline{R}$, $\overline{B}$, and $\overline{\Psi}$ are similarly defined.

It is clear from Lemma 1 that these ratios yield complex expressions that are not amenable to tractable analyses. To overcome this difficulty, we utilize the assumption introduced in §3.2 that noncompliance occurs rarely, i.e., $\lambda \ll \mu$. Expanding $\overline{I}$, $\overline{R}$, $\overline{B}$, and $\overline{\Psi}$ with respect to the ratio λ/μ and retaining up to the first-order terms yield the following:

Corollary 1 (Approximate measures under periodic inspections) *(i)* $\overline{I} = \frac{1}{\tau}\left(1 - \frac{\lambda}{\mu}\left(\frac{s}{\tau} + \frac{e^{-\mu s} - e^{-\mu\tau}}{\mu\tau}\right)\right) - \lambda\left(\frac{s}{\tau} - \frac{s^2}{2\tau^2}\right)$; *(ii)* $\overline{R} = \frac{\lambda}{\mu}\left(\frac{s}{\tau} + \frac{e^{-\mu s} - e^{-\mu\tau}}{\mu\tau}\right)$; *(iii)* $\overline{B} = \frac{\lambda}{\mu}\left(1 - \frac{s}{\tau} - \frac{e^{-\mu s} - e^{-\mu\tau}}{\mu\tau}\right)$; *(iv)* $\overline{\Psi} = \frac{\lambda}{\mu}\left(r\frac{s}{\tau} + (r + \kappa\mu)\frac{e^{-\mu s} - e^{-\mu\tau}}{\mu\tau}\right)$.

These approximations simplify the expressions substantially, enabling tractability. We use them as the basis of our analysis in the remainder of the paper. Finally, the long-run average social cost $\overline{C}$ is evaluated using the performance measures above via the relation $\overline{C} = \chi\overline{I} + r\overline{R} + h\overline{B}$ (see §3.5):

$$\overline{C} = \frac{h\lambda}{\mu} + \frac{\chi}{\tau} - \chi\lambda\left(\frac{s}{\tau} - \frac{s^2}{2\tau^2}\right) - \frac{\lambda}{\mu}\left(\frac{\chi}{\tau} + h - r\right)\left(\frac{s}{\tau} + \frac{e^{-\mu s} - e^{-\mu\tau}}{\mu\tau}\right). \tag{2}$$

4.1.2 Equilibrium

Recognizing that nondisclosure and full disclosure are obtained by setting $s = 0$ and $s = \tau$, we see that the objective functions under the binary disclosure policy follow directly from the expressions derived above. With the objective functions for both disclosure policies specified, we now characterize the equilibrium arising under periodic inspections. As a first step of the backward induction, we start with the firm's optimal response to the announced values of $\kappa \geq 0$ and $\tau > 0$.

Lemma 2 *Under periodic inspections with $\kappa \geq 0$ and $\tau > 0$, the firm chooses the threshold disclosure policy with* $s^* = \min\left\{\frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right), \tau\right\}$.

As expected, the firm prefers the threshold disclosure policy to the binary disclosure policy because the former presents flexibility that the latter lacks (continuous decision variable s vs. binary decision). Notice from the lemma that nondisclosure ($s^* = 0$) occurs if and only if $\kappa = 0$. That is, the firm keeps silence if no penalty is charged for noncompliance detection, but even a small amount of penalty prompts the firm to adopt partial disclosure ($s^* > 0$).

Next, we turn to the regulator's problem. Anticipating the firm's choice specified above, the regulator who employs periodic inspections determines κ and τ that minimize the long-run average social cost subject to the maximum penalty constraint. Let $\overline{I}^*$, $\overline{R}^*$, $\overline{B}^*$, and $\overline{\Psi}^*$ be the performance measures in Corollary 1 evaluated at s^*. Then the regulator's problem is to minimize $\overline{C}^* \equiv \chi\overline{I}^* +$

$r\overline{R}^* + h\overline{B}^*$ subject to $\kappa \leq r\delta$. For notational convenience, let

$$\sigma \equiv \frac{1}{\mu} \ln\left(1 + \mu\delta\right). \quad (3)$$

The equilibrium decisions, denoted by the superscript p, are as follows.

Proposition 1 (Equilibrium under periodic inspections) *Let* $G(\tau|\sigma) \equiv \left(1 + \frac{h-r}{\chi}\tau\right)(1 - e^{-\mu\tau}) - \left(2 + \frac{h-r}{\chi}\tau\right)\left(1 - \frac{\sigma}{\tau} - \frac{e^{-\mu\sigma} - e^{-\mu\tau}}{\mu\tau}\right) + \mu\sigma\left(1 - \frac{\sigma}{\tau}\right)$. *In equilibrium, the regulator employing the periodic inspection policy chooses* $\kappa^p = r\delta$ *and* $\tau^p = \max\{\sigma, \widehat{\tau}(\sigma)\}$, *where* $\widehat{\tau}(\sigma) > 0$ *is the unique solution to the equation* $G(\tau|\sigma) = \frac{\mu}{\lambda} - 1$. *In response, the firm chooses* $s^p = \sigma$.

A number of important observations are made from this proposition. First, the maximum penalty constraint binds in equilibrium, i.e., the regulator sets the penalty κ to the maximum allowed amount $r\delta$. This agrees with the results found in the majority of papers in the law enforcement literature including Becker (1968) and Kaplow and Shavell (1994), and it is in part driven by the assumption that the regulator may levy the penalty without incurring any cost of her own. Second, in equilibrium the firm's choice of the disclosure window size s^p is independent of the inspection interval τ^p, despite the fact that in general the firm's optimal response s^* is a function of both κ and τ (see Lemma 2). This happens because, once the regulator uses the penalty as a lever for inducing the firm to choose a desired size of the disclosure window, she adjusts the inspection interval to ensure that it is never exceeded by the window size (i.e., $s \leq \tau$ should be maintained); otherwise inspections are performed too frequently, not taking full advantage of the firm's disclosure. We discuss other implications of Proposition 1 in §5, where we compare the equilibria under periodic and random inspections.

4.2 Random Inspections

As described in §3.4, under the random inspection policy the regulator sets the inspection frequency $\nu = 1/\tau$ but randomizes the actual inspection times by sampling from the exponential distribution to determine the inspection interval T. As the next lemma shows, the firm responds to the regulator's announcement of the penalty and inspection frequency in a much simpler manner than he does under periodic inspections.

Lemma 3 *Under random inspections with* $\kappa \geq 0$ *and* $\tau > 0$, *the firm chooses the binary disclosure policy under which he does not disclose noncompliance if* $\kappa < r\tau$ *while he fully discloses noncompliance if* $\kappa \geq r\tau$. *The performance measures in each case are: (i)* $\overline{I}^* = \frac{\mu}{\lambda+\mu}\left(\frac{1}{\tau} + \frac{\lambda}{\mu\tau+1}\right)$, $\overline{R}^* = \frac{\lambda}{\lambda+\mu}\frac{1}{\mu\tau+1}$, $\overline{B}^* = \frac{\lambda}{\lambda+\mu}\frac{\mu\tau}{\mu\tau+1}$, *and* $\overline{\Psi}^* = \frac{\lambda}{\lambda+\mu}\frac{r+\kappa\mu}{\mu\tau+1}$ *if* $\kappa < r\tau$; *(ii)* $\overline{I}^* = \frac{\mu}{\lambda+\mu}\frac{1}{\tau}$, $\overline{R}^* = \frac{\lambda}{\lambda+\mu}$, $\overline{B}^* = 0$, *and* $\overline{\Psi}^* = \frac{r\lambda}{\lambda+\mu}$ *if* $\kappa \geq r\tau$.

In contrast to the periodic inspections case, the threshold disclosure policy does not offer an advantage over the binary disclosure policy to the firm. In fact, the optimal threshold policy under random inspections degenerates into the optimal binary policy. Such a simple decision structure is enabled by memorylessness of the inspection interval T. Because of memorylessness, the firm's expectation about the timing of the next inspection does not change over time. Then, due to this time symmetry, the decision on whether or not an instance of noncompliance should be disclosed applies uniformly to all occurrences of noncompliance: either all are disclosed or none is. As a result, partial disclosure never arises under random inspections.

Recall from §4.1 that we used approximate performance measures to derive the equilibrium solution under periodic inspections. To maintain consistency, we similarly approximate the measures in Lemma 3 and use them as the basis of our analysis:

Corollary 2 (Approximate measures under random inspections) *If $\kappa < r\tau$, then: (i) $\overline{I}^* = \frac{1}{\tau}\left(1 - \frac{\lambda}{\mu}\frac{1}{\mu\tau+1}\right)$; (ii) $\overline{R}^* = \frac{\lambda}{\mu}\frac{1}{\mu\tau+1}$; (iii) $\overline{B}^* = \frac{\lambda}{\mu}\frac{\mu\tau}{\mu\tau+1}$; (iv) $\overline{\Psi}^* = \frac{\lambda}{\mu}\frac{r+\kappa\mu}{\mu\tau+1}$. If $\kappa \geq r\tau$, then: (i) $\overline{I}^* = \frac{1}{\tau}\left(1 - \frac{\lambda}{\mu}\right)$; (ii) $\overline{R}^* = \frac{\lambda}{\mu}$; (iii) $\overline{B}^* = 0$; (iv) $\overline{\Psi}^* = \frac{r\lambda}{\mu}$.*

Using these expressions, we can write the long-run average social cost $\overline{C}^* = \chi\overline{I}^* + r\overline{R}^* + h\overline{B}^*$ as

$$\overline{C}^* = \begin{cases} \frac{r\lambda}{\mu} + \frac{\chi}{\tau}\left(1 - \frac{\lambda}{\mu}\right) & \text{if} \quad \tau \leq \frac{\kappa}{r}, \\ \frac{h\lambda}{\mu} + \frac{\chi}{\tau} - \frac{\lambda}{\mu}\left(\frac{\chi}{\tau} + h - r\right)\frac{1}{\mu\tau+1} & \text{if} \quad \tau > \frac{\kappa}{r}. \end{cases} \tag{4}$$

The regulator chooses κ and τ that together minimize this function subject to the maximum penalty constraint $\kappa \leq r\delta$. The equilibrium decisions, denoted by the superscript r, are as follows.

Proposition 2 (Equilibrium under random inspections) *Let $\tau^\dagger \equiv \left(2\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\frac{\lambda\mu}{\mu-\lambda}} + \frac{\lambda\mu}{\mu-\lambda} - \mu\right)^{-1}$ and $\tau^\ddagger \equiv \left(\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\frac{\lambda\mu}{\mu-\lambda}} - \mu\right)^{-1}$. In equilibrium, the regulator employing the random inspection policy chooses $\kappa^r \in [0, r\delta]$ and $\tau^r = \tau^\ddagger$ if $\delta < \tau^\dagger$ while she chooses $\kappa^r = r\delta$ and $\tau^r = \delta$ if $\delta \geq \tau^\dagger$. In response, the firm chooses nondisclosure if $\delta < \tau^\dagger$ while he chooses full disclosure if $\delta \geq \tau^\dagger$.*

Note that the quantities $\tau^\dagger$ and $\tau^\ddagger$ defined in the proposition satisfy $0 < \tau^\dagger < \tau^\ddagger$ under the assumptions $\chi < (h-r)\lambda/\mu^2$ and $\lambda \ll \mu$ introduced in §3. The proposition confirms the dichotomous structure of the equilibrium under random inspections which follows directly from the firm's binary response. In addition, it reveals a major difference between the equilibria under periodic and random inspections: under the latter, the maximum penalty constraint does not necessarily bind. For small values of δ, the regulator may choose any penalty amount κ without changing the equilibrium as long as it satisfies the constraint $\kappa \leq r\delta$. This degenerate solution arises if the regulator finds full disclosure inefficient. That is, if a large penalty cannot be imposed (small δ), even at the maximum allowed penalty the regulator may be unable to find a cost-effective inspection frequency that induces full disclosure. Facing such a situation, the regulator will instead choose the inspection frequency that minimizes the social cost without inducing full disclosure. When this equilibrium is established, however, an incremental increase in penalty does not change the firm's nondisclosure response nor the social cost. The regulator's indifference toward penalty for the case $\delta < \tau^\dagger$ reflects this situation.

Combining (4) with the results in Proposition 2, we can obtain a closed-form expression for the optimal long-run average social cost:

$$\overline{C}^r = \begin{cases} \frac{r\lambda}{\mu} - \chi(\mu - 2\lambda) + 2\chi\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\left(1 - \frac{\lambda}{\mu}\right)\lambda} & \text{if} \quad \delta < \tau^\dagger, \\ \frac{r\lambda}{\mu} + \frac{\chi}{\delta}\left(1 - \frac{\lambda}{\mu}\right) & \text{if} \quad \delta \geq \tau^\dagger. \end{cases} \tag{5}$$

5 Comparison of Equilibria

In this section we compare the equilibria under periodic and random inspections that we derived in the last section. In §5.1 we focus on the relationship between the two enforcement levers under each inspection policy, and in §5.2 we compare the social costs under the two policies. In the discussions below we pay special attention to how the equilibria are impacted by the maximum penalty constraint, represented by the parameter δ, which plays a key role in determining the effectiveness of an inspection policy.

5.1 Relationship between Penalty and Inspection Frequency

We first examine the relationship between the two enforcement levers, penalty and inspection frequency, when their values are determined in equilibrium. As found in Proposition 1 and Proposition 2, under both periodic and random inspections the penalties κ^p and κ^r are set to the maximum amount ($r\delta$) except when δ is sufficiently small under random inspections (see §4.2). That the maximum penalty limit is reached reflects the fact that the penalty is a more efficient instrument than inspections; unlike the latter, the regulator does not incur any direct cost by using the penalty as an incentive. It also implies that relaxing the binding maximum penalty constraint (larger δ) leads to a new equilibrium in which the penalty is increased accordingly. Thus, κ^p and κ^r increase as δ becomes larger.

By contrast, the behaviors of τ^p and τ^r—or equivalently the inspection frequencies—are not straightforward, as the next proposition reveals. (Note that $\widehat{\tau}(\sigma)$ and $\tau^\dagger$ appearing in the proposition are defined in Proposition 1 and Proposition 2.)

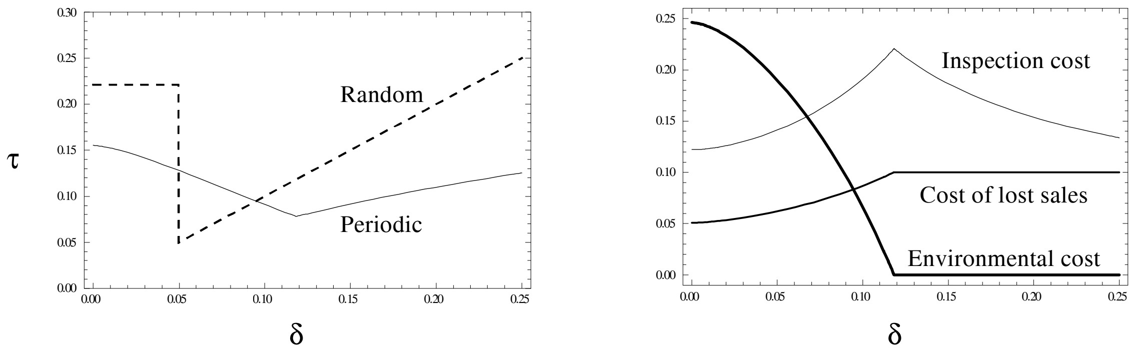

(a) Mean inspection intervals τ^p and τ^r

(b) Cost components under periodic inspections

Figure 3: The figure on the left shows the equilibrium mean inspection intervals τ^p (solid line) and τ^r (dashed line) as a function of the normalized maximum penalty δ. The figure on the right shows three components of the equilibrium long-run average social cost under periodic inspections. The following parameter values are used in this example: $\chi = 0.02$, $h = 5$, $r = 1$, $\lambda = 1$, and $\mu = 10$.

Proposition 3 *As δ increases from zero to infinity, τ^p initially decreases but increases after reaching a point at which $\sigma = \widehat{\tau}(\sigma)$. On the other hand, τ^r initially stays constant until it jumps downward at $\delta = \tau^\dagger$, after which it increases. Both τ^p and τ^r increase if and only if the firm chooses full disclosure.*

For sufficiently large δ, both τ^p and τ^r increase in δ. That is, the regulator performs less frequent inspections if she is able to charge higher penalty for a violation. This relationship is intuitive since higher penalty lessens the need to perform costly inspections—this is exactly the substitutability known in the literature. However, the proposition states that substitutability does not hold when δ is small. For small δ, the opposite may happen: under periodic inspections, the two enforcement levers in fact act as *complements*, i.e., higher penalty should be accompanied by more frequent inspections rather than lessen the need for inspections. See Figure 3(a) for an illustration of this non-monotonic behavior. As a direct consequence, the long-run average cost of inspections ($\chi\overline{I}$) also exhibits non-monotonicity; see Figure 3(b) that illustrates this for periodic inspections. From this finding, we conclude that the nature of the interaction between the two levers may fundamentally change depending on how much penalty can be levied.

Substitutability may be reversed under periodic inspections because transparency of the inspection schedule allows the firm to fine-tune his disclosure timing. If δ is small, such an opportunistic behavior may result in a partial disclosure equilibrium in which the firm does not disclose early occurrences of noncompliance. Consider the chain of events that unfolds after the penalty is increased incrementally, starting from a partial equilibrium. In order to avoid the increased penalty, the firm responds to this change by disclosing more early noncompliance occurrences (i.e., choose a larger disclosure window), thereby preempting inspections. This leads to fewer inspections, which benefit the regulator in the form of reduced cost of inspections. However, she does not simply absorb this saving. The regulator reinvests the saving in performing more inspections, because she can induce even more disclosure by doing so as long as full disclosure has not been reached. Hence, a reinforcement mechanism is in effect: unless the firm fully discloses noncompliance, an increased penalty leads to a new equilibrium in which the regulator schedules more frequent inspections.

The reversal of substitutability exists under random inspections as well, even though partial disclosure equilibrium does not arise. This is evident from Figure 3(a), which shows that τ^r jumps to a lower number, i.e., the equilibrium inspection frequency under random inspections jumps to a higher number, as the firm switches from nondisclosure to full disclosure in response to the increased penalty. This is an extreme version of complementarity, pushed to the limit by the firm's binary disclosure decision. Thus the non-monotonic relationship between the two enforcement levers is quite general.

We therefore conclude that the substitutability between penalty and inspection intensity—the accepted notion in the probabilistic law enforcement literature—is in fact a qualified truth. If full disclosure cannot be induced because of the limited amount of penalty that can be charged to the firm, substitutability is replaced by complementarity.

5.2 Relative Social Costs and Choice of Inspection Policy

Next, we compare the equilibrium long-run average social costs under periodic and random inspections, denoted by $\overline{C}^p$ and $\overline{C}^r$. This comparison offers a useful guidance to the regulator who faces a choice between the two inspection policies. As before, we pay special attention to the maximum penalty constraint. To this end, we examine how the social costs vary with δ under the two inspection policies. The following proposition focuses on the limiting cases.

Proposition 4 *$\overline{C}^p < \overline{C}^r$ for sufficiently small δ, whereas $\overline{C}^p > \overline{C}^r$ for sufficiently large δ.*

In other words, random inspections are preferred if and only if the regulator can charge a large penalty to the firm (large δ); otherwise, periodic inspections are preferred. The numerical example presented in Figure 4(a) confirms this finding. As the figure shows, the advantage of random inspections disappears as δ approaches zero, the limit at which a lower social cost is attained under periodic inspections than under random inspections.

Intuition suggests that performing random inspections is more effective than performing periodic inspections. Consider periodic inspections. Under this policy the firm perfectly anticipates when the next inspection arrives, and therefore he is able to fine-tune his disclosure timing so that he discloses only the noncompliance episodes that occur close to the next inspection arrival. Such an opportunistic behavior is mitigated under the random inspection policy, since there is no certainty about the timing

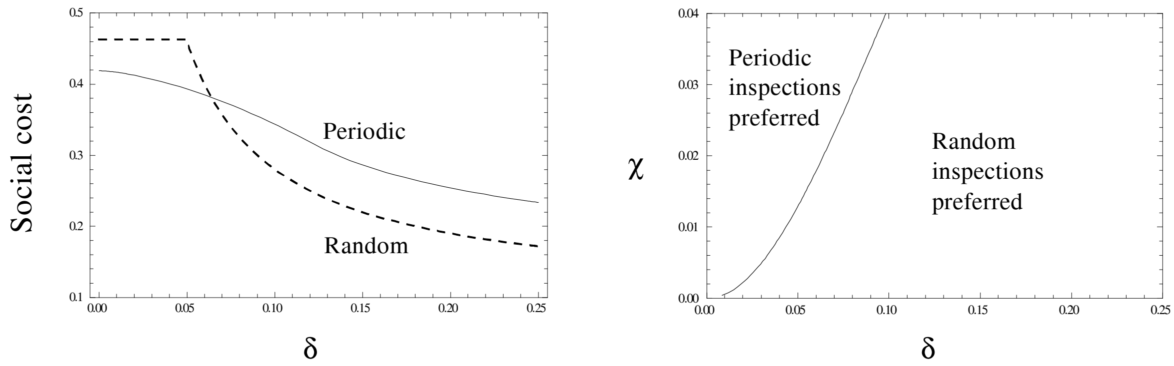

(a) Equilibrium social costs

(b) Inspection policy preference map

Figure 4: The figure on the left shows the equilibrium long-run average social costs $\overline{C}^p$ (solid line) and $\overline{C}^r$ (dashed line) as a function of the normalized maximum penalty δ. The figure on the right shows the regions in the (δ, χ) space in which one of the two inspection policies is preferred from the social cost perspective. The same parameter values as in Figure 3 are chosen, except that χ in the second figure is varied between 0 and 0.04.

of inspections; the "surprise factor" inherent in the random inspection policy would lower the firm's ability to plan ahead. According to this reasoning, then, the random inspection policy should dominate the periodic inspection policy. This is indeed what happens when δ is large, but Proposition 4 reveals that the opposite is true when δ is small: the random inspection policy is dominated by the periodic inspection policy. What drives this reversal?

As it turns out, *ceteris paribus*, it is more efficient to perform periodic inspections than to perform random inspections when the compliance/noncompliance states alternate as a Markov process. The next lemma supports this assertion.

Lemma 4 *With fixed inspection frequency, the long-run average social cost is lower under periodic inspections than under random inspections if the firm is induced to choose either nondisclosure or full disclosure under both inspection policies.*

Note that the lemma is proved using the exact performance measures; the result is general and does not depend on the approximations. The assumptions in Lemma 4 are chosen to isolate the impact of randomization, the key difference between the two inspection policies. As a result, unlike in Proposition 4, we do not compare the equilibrium outcomes that are influenced by other confounding factors such as the binding maximum penalty constraint. Instead, we control for these factors as follows. First, we normalize the firm's disclosure behavior by considering only the cases in which the firm either fully discloses noncompliance or never does under both inspection policies. (Partial disclosure is not considered because such an outcome does not arise under random inspections.) Second, we fix the inspection frequency, which allows us to focus on the effect of randomizing the inspection interval around its mean. With these factors controlled for, the lemma states that performing periodic inspections is more efficient than performing random inspections.

To understand the reason for this result, consider the case where the firm never discloses noncompliance. Then noncompliance can only be discovered via detection, and therefore the efficiency of an inspection policy is determined entirely by the detection probability. This is precisely $\theta(t)$ given in (1), namely, the probability that an inspection at time t finds noncompliance after compliance is observed at time zero. Under the periodic inspection policy, inspections are performed every constant time units equal to τ; hence, the detection probability is $\theta(\tau)$. On the other hand, under the random inspection policy, inspections are performed every T random time units, which has mean $E[T] = \tau$; hence the detection probability is $E[\theta(T)]$ in this case. Notice from (1) that $\theta(t)$ is concave increasing. Then by Jensen's inequality we have $\theta(\tau) > E[\theta(T)]$, i.e., the probability of detection is greater under periodic inspections than under random inspections. This implies that periodic inspections are more efficient than random inspections at detecting noncompliance, thus lowering the overall cost.

Therefore, higher efficiency is achieved with periodic inspections because the detection probability exhibits diminishing returns in time. This concavity arises from the transient behavior of the underlying Markov process, which restarts each time an inspection arrives and finds compliance (due to memorylessness of the compliance duration). Once compliance is observed, it takes time for the ensuing state transitions to gradually settle into the steady state. Then, as the regulator delays the time of next inspection, the probability of finding noncompliance at that time increases but with decreasing rates because the probability converges to the steady-state limit $\lim_{t\to\infty} \theta(t) = \lambda/(\lambda + \mu)$. This convergence gives rise to concavity, which is significant as long as inspections are performed at finite intervals.[8]

The effect of concavity is as follows. Consider a regulator who flips a coin to decide whether she should inspect the firm earlier or later than the scheduled inspection time. An early detection presents an advantage of preventing more pollutions, but because the detection probability increases in time, the chance of detection is higher if the inspection is delayed. Although randomization balances this tradeoff between early and late inspections, it also brings efficiency loss; due to the concavity in detection probability, randomization assigns more weight to an early non-detection than to a late detection. As a result, on average it leads to a lower probability of detection.

Returning to the result of Proposition 4, we now see that the conditional statement there—that random inspections are preferred if and only if δ is large—originates from the tension between efficiency of noncompliance detection and the firm's opportunistic disclosure behavior. Relative to random inspections, performing periodic inspections brings higher efficiency of detection but at the same time allows the firm to fine-tune his disclosure timing, thus softening the impact of an increased penalty. When δ is small, the former benefit dominates the latter disadvantage because the impact of the firm's muted response to the penalty is limited by the small amount charged while the relative contribution of inspections is increased. When δ is large, by contrast, the opposite happens because the firm's muted response is amplified by the large penalty while the relative contribution of inspections is reduced.

Because of this tradeoff, a regulator who considers implementing a cost-effective inspection policy would have to make a choice. One key determinant of such a choice is the (normalized) maximum penalty δ, which has been our focus so far. Another is χ, the cost incurred in each inspection. Both are important economic factors that directly influence the effectiveness of an inspection policy. In Figure 4(b) we identify the regions in the (δ, χ) space in which either the periodic inspection policy or the random inspection policy is preferred in terms of minimizing the social cost. As the figure illustrates, periodic inspections are preferred if the amount of penalty that can be charged to the firm is limited and it is expensive to perform an inspection. Otherwise, random inspections are preferred. Therefore, periodic inspections are recommended to a regulator who operates in restrictive conditions (small δ and large χ). In the next section we discuss the implications of this finding.

6 Conclusions

In this paper we offer new perspectives on the problem of environmental regulation enforcements. Building on the ideas and tools from the mathematical theory of reliability, we develop a novel analytical framework that brings precision to the problem features that have been abstracted in the existing literature. In particular, we add a time dimension to the decisions made by a manufacturing firm and a regulator in a production setting, representing stochastic pollutant emissions as an alternating Markov process. We study the situations in which environmental violations occur unintentionally. We assume that the firm follows a decision rule on when such random occurrences should be disclosed, knowing that he will be charged a penalty if noncompliance is discovered in an inspection. In turn, the regulator chooses the amount of penalty and the frequency of inspections that minimize the long-run average social cost. The regulator employs either periodic inspections performed at regular intervals or random inspections, under which the inspection interval is sampled from an exponential distribution.

We find that, contrary to a commonly held belief, the two enforcement levers—inspection intensity and penalty—do not necessarily act as substitutes in influencing the firm's decision. While it is intuitive that a higher penalty alleviates the need for frequent inspections, such a relationship is valid only when the firm is induced to fully disclose noncompliance. Full disclosure may not happen, though. If the maximum amount of penalty that can be charged to the firm is limited, the firm either never discloses noncompliance (under random inspections) or discloses only the late occurrences of noncompliance (under periodic inspections), taking a chance that the early occurrences will be resolved before the next scheduled inspection. With such an opportunistic behavior, substitutability may be reversed: the two enforcement levers act as complements. This means that the regulator minimizes the social cost by combining frequent inspections with a larger penalty, thus offering a stick-and-stick rather than a carrot-and-stick incentive to the firm. It also implies that the regulator's enforcement cost may rise despite the threat of increased penalty, the opposite of the known conclusion in the literature.

Furthermore, we find that periodic inspections may outperform random inspections even though the latter possesses a "surprise factor" that the former lacks. This is not immediately clear, because periodic inspections provide the firm with perfect information about the inspection schedule, which the firm uses to act opportunistically with regard to his disclosure timing. Such a behavior lowers the effectiveness of periodic inspections, yet the net efficiency may actually be higher than under random inspections. As it turns out, periodic inspections are more efficient at detecting noncompliance when compliance/noncompliance states alternate stochastically. This happens because finite inspection intervals capture the transient behavior of the state transitions, which introduces efficiency loss when inspections are randomized. Therefore, there exists a tradeoff between efficient noncompliance detection and exacerbating the firm's opportunistic disclosure behavior. Depending on which is more significant, either inspection policy may be preferred. We find that periodic inspections are preferred when the regulator operates in restrictive conditions, marked by limited penalty and high cost of inspections. If these conditions are not in place, random inspections are preferred.

Our findings suggest that the strategy for enforcing environmental regulations should be tailored to the characteristics of a firm. If a significant portion of environmental violations occur unintentionally due to accidents or other random causes such as equipment malfunctions, a one-size-fits-all approach is inadequate. Against large firms with sufficient financial means and transparent operations (e.g., multinational manufacturers), the strategy of utilizing random inspections and imposing large fines in lieu of inspections would be effective. However, according to our analysis, the same strategy would not be as effective against small firms with limited financial resources and nontransparent operations (e.g., local suppliers). For these firms, the regulator should consider complementing penalties with frequent inspections, performed at regular intervals. Recognizing this difference is especially important because in many cases the worst polluters are small-scale enterprises operating in less visible sectors. As Lanjouw (2006, p. 54) notes, it is believed that "small firms are more intensive producers of pollution than large firms" in large part because these firms tend to use older equipment that are more prone to breakdowns. Given the potentially large environmental impact originating these firms, a well-executed enforcement strategy tailored to them may have a disproportionate effect on overall performance of environmental regulations.

Finally, the analytical framework developed in this paper has a potential as a diagnostic tool. For example, the framework offers a natural way to infer the amount of emissions that escape detection, which can be readily computed by evaluating the performance measures derived in the model. Such inferences are widespread in other areas (for example, Kaplan (2010) suggests a queuing model that predicts the number of undetected terror plots), and our model paves the way for similar applications.

Notes

[1]In the literature the terms "audit" and "monitoring" are often used synonymously with "inspection," despite the differences in nuance. We adopt the same convention but primarily use the last term.

[2]The stochastic process of alternating compliance/noncompliance states is best understood using a machine repair analogy. Suppose that the firm operates a pollution control system that removes toxicity from the pollutant. Because of imperfect reliability the system goes down occasionally, emitting untreated pollutant unless production is suspended. Noncompliance due to the failed system lasts until system repair is completed. In this example the constant rates of state transitions correspond to system failure rate and average speed of repair. This example operationalizes the suggestion by Malik (1993) that equipment malfunction is one the major sources of stochastic emissions.

[3]The latter assumption is reasonable because disclosure will include information about the exact scope and location of a violation, directing the regulator to where it has occurred and sparing her effort to find out the details.

[4]Together with the earlier assumption $\lambda \ll \mu$ this condition implies $\chi \ll (h-r)/\mu$, which states that the inspection cost is negligible compared to the net cost incurred by the society while a noncompliance episode goes unreported (the expected duration of which is equal to $1/\mu$).

[5]If we take the view that the maximum penalty is determined by the firm's wealth, we may interpret δ and $r\delta$ as days-cash-on-hand and cash reserve, respectively.

[6]For example, setting $\mu/\lambda = 10$ in the given equation yields $\widehat{y} = 6.57 \times 10^7$, implying that the maximum penalty that can be levied should satisfy $\kappa = r\delta < r\widehat{y}/\mu = 0.1 \times \widehat{y} \times (r/\lambda)$, i.e., the penalty should be smaller than 6.57 million times the revenue earned for the expected duration of compliance. Increasing the ratio μ/λ further increases $\widehat{y}$ and hence the allowed penalty amount.

[7]Social welfare maximization criterion is ubiquitous in the regulation context, and most of the economics papers we reviewed in §2.2 assume the same objective. The objectives of maximizing or minimizing a long-run average performance measure are common in many OM models, including those for inventory controls such as the (Q, r) model (e.g., Zheng 1992). The games with static decision

variables in infinite horizon settings are also commonly found in the literature (e.g., Cachon and Zipkin 1999).

[8]Lemma 4 also shows that the periodic inspection policy continues to outperform the random inspection policy even when the firm fully discloses noncompliance. The reason is essentially the same as the case of nondisclosure. While no detection occurs in this case, higher probability of noncompliance detection is translated into a smaller expected number of inspections needed while compliance is maintained. Therefore, the long-run average inspection cost is lower under periodic inspections than under random inspections.

Appendix

A Proofs of Main Results

Proof of Lemma 1. See Appendix B. ■

Proof of Lemma 2. First, suppose that the firm chooses the threshold disclosure policy. Recall that the firm's choice s^* is defined in the interval $[0, \tau]$. Differentiating $\overline{\Psi}(s)$ in Corollary 1, $\overline{\Psi}'(s) = \frac{\lambda}{\mu\tau}\left(r - (r + \kappa\mu)\, e^{-\mu s}\right)$ and $\overline{\Psi}''(s) = \frac{\lambda}{\tau}(r + \kappa\mu)\, e^{-\mu s} > 0$. Hence, $\overline{\Psi}(s)$ is convex. If $\kappa = 0$ then $\overline{\Psi}'(0) = 0$ and therefore $\overline{\Psi}(s)$ is minimized at $s = 0$. Suppose $\kappa > 0$, in which case $\overline{\Psi}'(0) < 0$. If $\tau \leq \frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right)$ then $\lim_{s\to\tau}\overline{\Psi}'(s) = \frac{\lambda}{\mu\tau}\left(r - (r + \kappa\mu)\, e^{-\mu\tau}\right) \leq 0$ so $\overline{\Psi}(s)$ is minimized at the boundary $s = \tau$. Otherwise $\lim_{s\to\tau}\overline{\Psi}'(s) > 0$ so $\overline{\Psi}(s)$ is minimized at a unique interior point $s^* < \tau$ whose solution is found from the first-order condition $\overline{\Psi}'(s) = 0$, which yields $s = \frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right)$. Combining all cases, we have $s^* = \min\left\{\frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right), \tau\right\}$. Evaluating $\overline{\Psi}(s)$ at s^*, we get the reduced long-run average cost $\overline{\Psi}_\theta(\kappa)$, a continuous function defined as $\overline{\Psi}_\theta(\kappa) = \frac{r\lambda}{\mu^2\tau}\left(1 + \ln\left(1 + \frac{\kappa\mu}{r}\right) - \left(1 + \frac{\kappa\mu}{r}\right)e^{-\mu\tau}\right)$ for $\kappa < \kappa_\theta$ and $\overline{\Psi}_\theta(\kappa) = \frac{r\lambda}{\mu}$ for $\kappa \geq \kappa_\theta$ where $\kappa_\theta \equiv \frac{r}{\mu}(e^{\mu\tau} - 1)$. Next, suppose that the firm chooses the binary disclosure policy. Substituting the nondisclosure and full disclosure conditions $s = 0$ and $s = \tau$ in the expression for $\overline{\Psi}$ in Corollary 1 and comparing them, we find that the firm chooses full disclosure if and only if $\kappa > \kappa_b$ where $\kappa_b \equiv \frac{r}{\mu}\left(\frac{\mu\tau}{1 - e^{-\mu\tau}} - 1\right)$. The corresponding reduced long-run average cost is $\overline{\Psi}_b(\kappa)$, a continuous function defined as $\overline{\Psi}_b(\kappa) = \frac{\lambda}{\mu}(r + \kappa\mu)\frac{1 - e^{-\mu\tau}}{\mu\tau}$ for $\kappa \leq \kappa_b$ and $\overline{\Psi}_b(\kappa) = \frac{r\lambda}{\mu}$ for $\kappa > \kappa_b$. Now we compare $\overline{\Psi}_b(\kappa)$ with $\overline{\Psi}_\theta(\kappa)$. The following can be verified: (i) $\kappa_b < \kappa_\theta$; (ii) $\overline{\Psi}_\theta(0) = \overline{\Psi}_b(0)$; (iii) $\overline{\Psi}_b(\kappa)$ increases until $\kappa = \kappa_b < \kappa_\theta$ and then stays constant at $\overline{\Psi}_b(\kappa) = \frac{r\lambda}{\mu}$ afterwards; (iv) $\overline{\Psi}_\theta(\kappa)$ increases until $\kappa = \kappa_\theta > \kappa_b$ and then stays constant at $\overline{\Psi}_\theta(\kappa) = \frac{r\lambda}{\mu}$ afterwards; (v) $\overline{\Psi}_b'(\kappa) > \overline{\Psi}_\theta'(\kappa)$ for $\kappa \in (0, \kappa_b]$. Combined, these conditions imply $\overline{\Psi}_b(\kappa) \geq \overline{\Psi}_\theta(\kappa)$ for all $\kappa \geq 0$. Hence, the threshold disclosure policy weakly dominates the binary disclosure policy; the firm chooses the former. ■

Proof of Proposition 1. For notational convenience let us suppress the argument σ in $G(\tau|\sigma)$.

It can be proved that the equilibrium exists in the region $\tau \geq \sigma$ satisfying $\kappa = r\delta$, and that the firm sets $s^* = \sigma$ in equilibrium. (Proof is found in the accompanying Technical Appendix.) It then remains to find $\tau \geq \sigma$ that minimizes the regulator's long-run average social cost (2) with $s^* = \sigma$ substituted. Let $\overline{C}^*(\tau)$ be the reduced cost function with $s^* = \sigma$. First, we show that $\overline{C}^*(\tau)$ defined in the expanded region $\tau > 0$ has a unique interior minimizer. Differentiating $\overline{C}^*(\tau)$ and setting it to zero yields the first-order condition $G(\tau) = \frac{\mu}{\lambda} - 1$. Let us rewrite $G(\tau)$ as $G(\tau) = \left(1 + \frac{h-r}{\chi}\tau\right)(1 - e^{-\mu\tau}) - \left(2 + \frac{h-r}{\chi}\tau\right)\left(1 - \frac{1-e^{-\mu\tau}}{\mu\tau}\right) + \left(a - \frac{b}{\mu\tau}\right)$, where $a \equiv \mu\sigma + \frac{h-r}{\chi\mu}(\mu\sigma - 1 + e^{-\mu\sigma})$ and $b \equiv (1-\mu\sigma)^2 + 1 - 2e^{-\mu\sigma}$. It can be proved that $a = b = 0$ for $\sigma = 0$ while $a > 0$ and $b > 0$ for $\sigma > 0$. Differentiating $G(\tau)$ yields $G'(\tau) = \left(1 + \frac{h-r}{\chi}\tau\right)\mu e^{-\mu\tau} + \frac{2}{\tau}\left(e^{-\mu\tau} - \frac{1-e^{-\mu\tau}}{\mu\tau}\right) + \frac{b}{\mu\tau^2}$ and $G''(\tau) = \left(\frac{h-r}{\chi\mu}(1-\mu\tau) - \frac{2}{\mu\tau} - 1\right)\mu^2 e^{-\mu\tau} - \frac{4}{\tau^2}\left(e^{-\mu\tau} - \frac{1-e^{-\mu\tau}}{\mu\tau}\right) - \frac{2b}{\mu\tau^3}$. Moreover, the following properties hold: (i) in the limit $\tau \to 0$, $G(\tau)$, $G'(\tau)$, and $G''(\tau)$ approach 0, 0, and $\frac{(h-r)\mu}{\chi} - \frac{\mu^2}{3}$ if $\sigma = 0$, while they approach $-\infty$, ∞ and $-\infty$ if $\sigma > 0$; (ii) in the limit $\tau \to \infty$, $G(\tau)$ and $G'(\tau)$ approach $\frac{h-r}{\chi\mu} - 1 + a$ and 0. Since $G(\tau)$ exhibits behaviors that are qualitatively different around $\tau = 0$ depending on whether $\sigma = 0$ or $\sigma > 0$, we consider these two cases separately. Assume $\sigma = 0$ and let $G_0(\tau)$ be $G(\tau)$ with $\sigma = 0$. The limits shown above imply that $G_0(\tau)$ initially increases from zero at $\tau = 0$ if $\frac{h-r}{\chi\mu} > \frac{1}{3}$, converging to $\frac{h-r}{\chi\mu} - 1$ as $\tau \to \infty$. Let τ^0 be the solution of $G_0'(\tau) = 0$, i.e., τ^0 is a critical point of $G_0(\tau)$. Note that $G_0'(\tau^0) = 0$ can be written as $\varphi_0(\mu\tau^0) = \frac{h-r}{\chi\mu}(3 - \mu\tau^0)$, where $\varphi_0(x) \equiv \frac{1}{x^3}\left(2(e^x - 1 - x) - x^2\right)(3 - x)$. It can be proved that $\varphi_0(x) < 1$ for $x > 0$. (Proof is in the Technical Appendix.) Then, evaluating $G_0''(\tau)$ at τ^0, we get $G_0''(\tau^0) = \left(\varphi(\mu\tau^0) - 1\right)\mu^2 e^{-\mu\tau^0} < 0$. That $G_0''(\tau) < 0$ at a critical point τ^0 implies $G_0(\tau)$ is quasiconcave. Recall the assumptions $\frac{\mu}{\lambda} \gg 1$ and $\frac{h-r}{\chi\mu} > \frac{\mu}{\lambda}$ stated in §3 which together imply $\frac{h-r}{\chi\mu} \gg 1 > \frac{1}{3}$. Combined with the earlier finding that $G_0(\tau)$ initially increases if $\frac{h-r}{\chi\mu} > \frac{1}{3}$ and that $\lim_{\tau\to\infty} G_0(\tau) = \frac{h-r}{\chi\mu} - 1 > \frac{\mu}{\lambda} - 1$, quasiconcavity and continuity of $G_0(\tau)$ imply that $G_0(\tau)$ crosses $\frac{\mu}{\lambda} - 1$ exactly once from below at $\tau^0 > 0$. Therefore, the optimal solution is unique when $\sigma = 0$. Now assume $\sigma > 0$. Consider two cases: $b \geq 2$ and $0 < b < 2$. If $b \geq 2$, it is straightforward to show that $G'(\tau) > 0$ for all $\tau > 0$. Therefore, in this case $G(\tau)$ monotonically increases from $-\infty$ to $\frac{h-r}{\chi\mu} - 1 + a$ as τ goes from zero to ∞. Since $a > 0$ when $\sigma > 0$ and $\frac{h-r}{\chi\mu} > \frac{\mu}{\lambda}$, the limit $\frac{h-r}{\chi\mu} - 1 + a$ is greater than $\frac{\mu}{\lambda} - 1$; hence, $G(\tau)$ crosses $\frac{\mu}{\lambda} - 1$ exactly once from below and therefore the solution is unique. Now assume $b < 2$. In this case $G(\tau)$ may not be monotonically increasing, i.e., it may peak before it converges to $\frac{h-r}{\chi\mu} - 1 + a$ as $\tau \to \infty$. Let τ^σ be the solution of $G'(\tau) = 0$, i.e., τ^σ is a critical point of $G(\tau)$. Note that $G'(\tau^\sigma) = 0$ can be written as $\varphi(\mu\tau^\sigma) = \frac{h-r}{\chi\mu}(3 - \mu\tau^\sigma)$, where $\varphi(x) \equiv \frac{1}{x^3}\left((2-b)e^x - 2 - 2x - x^2\right)(3 - x)$. (The function $\varphi_0(x)$ defined above is a special case of $\varphi(x)$ with b set to zero.) It can be proved that $\varphi(x) < 1$ for $x > 0$. Then the argument similar to the case of $\sigma = 0$ leads to the conclusion that the solution is unique.

In all three cases ($\sigma = 0$, $\sigma > 0$ with $b \geq 2$, $\sigma > 0$ with $0 < b < 2$), we found that the equation $G(\tau) = \frac{\mu}{\lambda} - 1$ has a unique solution for which $G(\tau) < \frac{\mu}{\lambda} - 1$ to its left and $G(\tau) > \frac{\mu}{\lambda} - 1$ to its right. Since the solution depends on σ we denote it as $\widehat{\tau}(\sigma)$, which satisfies $G'(\widehat{\tau}(\sigma)) > 0$. Thus, $\overline{C}^*(\tau)$ defined in the expanded region $\tau > 0$ has a unique interior minimizer $\widehat{\tau}(\sigma)$, the statement we set out to prove. Restricting the region to that in which the equilibrium exists ($\tau \geq \sigma$), we finally conclude that in equilibrium the regulator chooses $\tau^p = \max\{\sigma, \widehat{\tau}(\sigma)\}$. ■

Proof of Lemma 3. Recall that the time between two successive and uninterrupted inspections is $T \sim \exp(\nu)$ where $\nu = \frac{1}{\tau}$. Suppose that the firm employs the threshold disclosure policy. At the moment noncompliance occurs, irrespective of when it occurs, the firm facing random inspections expects that the next inspection will arrive τ time units later; due to memorylessness of T, the expectation does not change. Then the firm discloses this occurrence if and only if $s = \tau$; if $s < \tau$, regardless of the value of s, the onset of noncompliance is always outside of the disclosure window and therefore the firm does not disclose it. Because the same logic applies to all occurrences of noncompliance, it follows that the firm discloses either all or none depending on whether $s = \tau$ or $s < \tau$. Hence, the threshold disclosure policy under random inspections degenerates into the binary disclosure policy. Next, we derive the conditions under which nondisclosure (ND) or full disclosure (FD) is optimal. The most convenient unit of analysis is a single compliance-noncompliance cycle ("cycle"), which forms a renewal. Each cycle lasts $U + D$ time units, where $U \sim \exp(\lambda)$ and $D \sim \exp(\mu)$. Without loss of generality, focus on a randomly selected cycle and set the cycle start to time zero. Let T_r be the remaining time until the next inspection since time U, which marks the onset of noncompliance. Due to memorylessness, T and T_r are identically distributed: $T_r \sim \exp(\nu)$. Under FD, the firm's expected profit per cycle is $E[\Pi] = rE[U] = \frac{r}{\lambda}$ since production lasts only up until noncompliance occurs. Under ND, on the other hand, the firm's expected profit per cycle is $E[\Pi] = rE[U] + rE[\min\{T_r, D\}] - \kappa \Pr(T_r \leq D) = \frac{r}{\lambda} + \frac{r - \kappa\nu}{\mu + \nu}$, reflecting the possibility that noncompliance is detected before it concludes (the event $T_r \leq D$). Since the cycle length is $E[U + D] = \frac{1}{\lambda} + \frac{1}{\mu}$ under both ND and FD, the corresponding long-run average costs for the firm are: $\overline{\Psi}^* = r - \frac{E[\Pi]}{E[U]+E[D]} = \frac{r\lambda}{\lambda+\mu}$ under FD and $\overline{\Psi}^* = r - \frac{E[\Pi]}{E[U]+E[D]} = \frac{\lambda}{\lambda+\mu}\frac{r+\kappa\mu}{\mu\tau+1}$ under ND. Comparing the two expressions, we find that the firm (weakly) prefers FD if $\kappa \geq r\tau$. Next, we derive the performance measures under ND and FD. Consider FD. Since all noncompliance episodes are disclosed, there exists no unsuspended noncompliance: hence, the expected durations of suspended and unsuspended noncompliance in a cycle are, respectively, $E[R] = E[D] = \frac{1}{\mu}$ and $E[B] = 0$. Consequently, $\overline{R} = \frac{E[R]}{E[U]+E[D]} = \frac{1/\mu}{1/\lambda+1/\mu} = \frac{\lambda}{\lambda+\mu}$ and $\overline{B} = 0$. To compute $\overline{I}$, note that inspections are performed only for the duration of U in a cycle since, unde FD, inspections are suspended as soon as noncompliance occurs.

Hence, the expected number of inspections performed in a cycle is equal to $E[I] = \nu E[U] = \frac{1}{\lambda\tau}$ and it follows that $\overline{I} = \frac{E[I]}{E[U]+E[D]} = \frac{1/\lambda\tau}{1/\lambda+1/\mu} = \frac{\mu}{\lambda+\mu}\frac{1}{\tau}$. Now consider ND. As we noted above, the probability of detection in a cycle is $\Pr(T_r \le D)$. Then $E[R] = E[D]\Pr(T_r \le D) = \frac{1}{\mu}\frac{\nu}{\mu+\nu} = \frac{1}{\mu}\frac{1}{\mu\tau+1}$ and $E[B] = E[\min\{T_r, D\}] = \frac{1}{\mu+\nu} = \frac{\tau}{\mu\tau+1}$. Hence, $\overline{R} = \frac{E[R]}{E[U]+E[D]} = \frac{\lambda}{\lambda+\mu}\frac{1}{\mu\tau+1}$ and $\overline{B} = \frac{E[B]}{E[U]+E[D]} = \frac{\lambda}{\lambda+\mu}\frac{\mu\tau}{\mu\tau+1}$. To compute $\overline{I}$, recall from above that the expected number of inspections for the first U time units in a cycle is $\nu E[U] = \frac{1}{\lambda\tau}$. In the next D time units the regulator performs at most one inspection: zero in the event $T_r > D$ and one in the event $T_r \le D$. The latter is true because inspections are suspended following a detection (the event $T_r \le D$). Therefore, the expected number of inspections in a cycle is $E[I] = \nu E[U] + 1 \cdot \Pr(T_r \le D) = \frac{1}{\lambda\tau} + \frac{1}{\mu\tau+1}$. Then $\overline{I} = \frac{E[I]}{E[U]+E[D]} = \frac{1/(\lambda\tau)+1/(\mu\tau+1)}{1/\lambda+1/\mu} = \frac{\mu}{\lambda+\mu}\left(\frac{1}{\tau} + \frac{\lambda}{\mu\tau+1}\right)$.
■

Proof of Proposition 2. Fix κ. Let $\overline{C}^*_-$ and $\overline{C}^*_+$ denote $\overline{C}^*$ defined in (4) for the regions $\tau \le \frac{\kappa}{r}$ and $\tau > \frac{\kappa}{r}$, respectively. Evaluating $\overline{C}^*_-$ and $\overline{C}^*_+$ at the boundary $\tau = \frac{\kappa}{r}$ and subtracting them yield $\overline{C}^*_+ - \overline{C}^*_- = \frac{\lambda}{\mu}\left(\frac{\chi r}{\kappa} + h - r\right)\frac{\mu\kappa}{\mu\kappa+r} > 0$, implying that $\overline{C}^*$ is discontinuous at $\tau = \frac{\kappa}{r}$, jumping upwards as τ crosses $\frac{\kappa}{r}$ from left to right. Hence, $\overline{C}^*_- < \overline{C}^*_+$ in the vicinity of $\tau = \frac{\kappa}{r}$. Suppose $\tau \le \frac{\kappa}{r}$. From (4) we see $\frac{\partial \overline{C}^*_-}{\partial \tau} < 0$, which implies that $\overline{C}^*_-$ is minimized at the boundary $\tau = \frac{\kappa}{r}$. Now suppose $\tau > \frac{\kappa}{r}$. Differentiating $\overline{C}^*_+$ and setting it to zero yields the solution $\tau^\ddagger = \left(\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\frac{\lambda\mu}{\mu-\lambda}} - \mu\right)^{-1}$. Substituting $\tau^\ddagger$ in the second derivative, we get $\left.\frac{\partial^2 \overline{C}^*_+}{\partial \tau^2}\right|_{\tau=\tau^\ddagger} = \frac{2\chi}{\tau^3}\left(1 - \frac{\lambda}{\mu}\right)\frac{1}{\mu\tau+1} > 0$, which implies that the unique critical point $\tau^\ddagger$ is a local minimizer of $\overline{C}^*_+$. It can be verified that $\lim_{\tau\to 0}\overline{C}^*_+ = \infty$, $\lim_{\tau\to\infty}\overline{C}^*_+ = \frac{h\lambda}{\mu}$, $\lim_{\tau\to 0}\frac{\partial \overline{C}^*_+}{\partial \tau} = -\infty$, and $\lim_{\tau\to\infty}\frac{\partial \overline{C}^*_+}{\partial \tau} = 0$; hence, $\tau^\ddagger$ is the unique interior minimizer of $\overline{C}^*_+$ in the expanded region $\tau \in (0, \infty)$. Whether $\tau^\ddagger$ is also the global minimizer depends on the conditions $\tau^\ddagger \le \frac{\kappa}{r}$ and $\tau^\ddagger > \frac{\kappa}{r}$. In the former case $\overline{C}^*_+$, defined in the region $\tau > \frac{\kappa}{r}$, is increasing in τ; together with the facts that $\overline{C}^*_-$ is minimized at $\tau = \frac{\kappa}{r}$ and that $\overline{C}^*_- < \overline{C}^*_+$ in the vicinity of $\tau = \frac{\kappa}{r}$, this implies that $\overline{C}^*$ is minimized at $\tau = \frac{\kappa}{r}$. In the latter case, $\overline{C}^*_+$ is minimized at the interior point $\tau^\ddagger > \frac{\kappa}{r}$ while $\overline{C}^*_-$ is minimized at $\tau = \frac{\kappa}{r}$. Given that $\overline{C}^*_- < \overline{C}^*_+$ in the vicinity of $\tau = \frac{\kappa}{r}$, there are two candidates for the global minimizer: $\tau = \frac{\kappa}{r}$ and $\tau = \tau^\ddagger$. Evaluating $\overline{C}^*$ at these values, we get $\overline{C}^*_1 = \frac{r\lambda}{\mu} + \frac{\chi r}{\kappa}\left(1 - \frac{\lambda}{\mu}\right)$ at $\tau = \frac{\kappa}{r}$ and $\overline{C}^*_2 = \frac{r\lambda}{\mu} - \chi(\mu - 2\lambda) + 2\chi\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\left(1 - \frac{\lambda}{\mu}\right)\lambda}$ at $\tau = \tau^\ddagger$. Then $\tau^\ddagger$ is the global minimizer if and only if $\overline{C}^*_2 < \overline{C}^*_1$, which is equivalent to the condition $\kappa < r\tau^\dagger$ where $\tau^\dagger = \left(2\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\frac{\lambda\mu}{\mu-\lambda}} + \frac{\lambda\mu}{\mu-\lambda} - \mu\right)^{-1}$. Note that, since $\tau^\dagger < \tau^\ddagger$, this condition implies $\kappa < r\tau^\ddagger$ or $\tau^\ddagger > \frac{\kappa}{r}$, which ensures the local minimum at $\tau = \tau^\ddagger$. Summarizing, the global minimizer of $\overline{C}^*$ is found at $\tau = \tau^\ddagger$ if $\kappa < r\tau^\dagger$ while it is at $\tau = \frac{\kappa}{r}$ if $\kappa \ge r\tau^\dagger$. The corresponding reduced social cost is $\overline{C}^{**} = \frac{r\lambda}{\mu} - \chi(\mu - 2\lambda) + 2\chi\sqrt{\left(\frac{h-r}{\chi} - \mu\right)\left(1 - \frac{\lambda}{\mu}\right)\lambda}$ if $\kappa < r\tau^\dagger$ and $\overline{C}^{**} = \frac{r\lambda}{\mu} + \frac{\chi r}{\kappa}\left(1 - \frac{\lambda}{\mu}\right)$ if $\kappa \ge r\tau^\dagger$. Notice that $\overline{C}^{**}$ for $\kappa < r\tau^\dagger$ is independent of κ whereas for $\kappa \ge r\tau^\dagger$ it decreases in κ. Hence, as κ

increases from zero to infinity, $\overline{C}^{**}$ stays constant for $\kappa < r\tau^{\dagger}$ and then decreases afterwards. Next, we incorporate the maximum penalty constraint $\kappa \leq r\delta$ in finding κ that minimizes $\overline{C}^{**}$. Two cases need to be considered separately: $\delta < \tau^{\dagger}$ and $\delta \geq \tau^{\dagger}$. If $\delta < \tau^{\dagger}$, the maximum penalty constraint implies $\kappa < r\tau^{\dagger}$, the region in which $\overline{C}^{**}$ does not vary with κ. Hence, the minimum is found at any value of κ satisfying $\kappa \leq r\delta$. Recall from above that the optimal τ in this case is $\tau = \tau^{\ddagger}$ since $\kappa < r\tau^{\dagger}$. If $\delta \geq \tau^{\dagger}$, on the other hand, the constraint $\kappa \leq r\delta$ includes the region in which $\overline{C}^{**}$ decreases in κ. Hence, $\overline{C}^{**}$ is minimized at the constraint boundary $\kappa = r\delta$. Recall from above that the optimal τ in this case is $\tau = \frac{\kappa}{r} = \delta$ since $\kappa = r\delta \geq r\tau^{\dagger}$. Summarizing, it is optimal for the regulator to choose $\kappa \in [0, r\delta]$ and $\tau = \tau^{\ddagger}$ if $\delta < \tau^{\dagger}$ and to choose $\kappa = r\delta$ and $\tau = \delta$ if $\delta \geq \tau^{\dagger}$. The firm's optimal response to this choice follows directly from Lemma 3. ■

Proof of Proposition 3. First consider periodic inspections. Recall from Proposition 1 that $\tau^p = \max\{\sigma, \widehat{\tau}\}$, where we suppressed the argument in $\widehat{\tau}(\sigma)$ for notational convenience. Suppose $\sigma < \widehat{\tau}$. Then $\tau^p = \widehat{\tau}$, which satisfies $G(\widehat{\tau}|\sigma) = \frac{\mu}{\lambda} - 1$ (Proposition 1). Implicit differentiation of this equation yields $\frac{d\widehat{\tau}}{d\sigma} = -\left.\frac{\partial G(\tau|\sigma)/\partial\sigma}{\partial G(\tau|\sigma)/\partial\tau}\right|_{\tau=\widehat{\tau}}$, which we now show to be negative. In the proof of Proposition 1 we showed that the numerator satisfies $\left.\frac{\partial G(\tau|\sigma)}{\partial\tau}\right|_{\tau=\widehat{\tau}} > 0$. Next, observe that the denominator satisfies $\left.\frac{\partial G(\tau|\sigma)}{\partial\sigma}\right|_{\tau=\widehat{\tau}} = \mu\left(1 - \frac{2\sigma}{\widehat{\tau}}\right) + \left(2 + \frac{h-r}{\chi}\widehat{\tau}\right)\frac{1-e^{-\mu\sigma}}{\widehat{\tau}} > \mu + \mu\left(\frac{h-r}{\chi\mu} - \frac{2\sigma}{\widehat{\tau}}\right)\frac{\mu\sigma}{1+\mu\sigma} > \mu + \mu\left(\frac{h-r}{\chi\mu} - 2\right)\frac{\mu\sigma}{1+\mu\sigma} > 0$, where the first inequality follows from $1 - e^{-x} > \frac{x}{1+x}$, the second inequality from the assumption $\sigma < \widehat{\tau}$ stated above, and the third from the assumptions $\frac{h-r}{\chi\mu} > \frac{\mu}{\lambda}$ and $\frac{\mu}{\lambda} \gg 1$ stated in §3 which together imply $\frac{h-r}{\chi\mu} > 2$. Therefore, $\frac{d\widehat{\tau}}{d\sigma} < 0$, i.e., $\widehat{\tau}$ decreases in σ if $\sigma < \widehat{\tau}$. Next, suppose $\sigma > \widehat{\tau}$. Then $\tau^p = \sigma$, which increases in σ. In summary, τ^p decreases in σ for $\sigma < \widehat{\tau}$ while it increases in σ for $\sigma > \widehat{\tau}$, reaching the minimum at $\sigma = \widehat{\tau}$. According to Proposition 1 full disclosure occurs if and only if $\sigma \geq \widehat{\tau}$, when $s^p = \tau^p = \sigma$. Since σ defined in (3) is an increasing function of δ, the statements in the proposition about τ^p follow. Now consider random inspections. Proposition 2 states that $\tau^r = \tau^{\ddagger}$ if $\delta < \tau^{\dagger}$ while $\tau^r = \delta$ if $\delta \geq \tau^{\dagger}$, where $\tau^{\ddagger}$ is independent of δ. Hence, τ^r does not vary in δ for $\delta < \tau^{\dagger}$ while it increases in δ for $\delta \geq \tau^{\dagger}$. From the same proposition we also see that the latter condition is sufficient and necessary for full disclosure. Note also that $\tau^{\dagger} < \tau^{\ddagger}$. Then τ^r jumps downward at $\delta = \tau^{\dagger}$ because, at that boundary, $\lim_{\delta\to(\tau^{\dagger})^-} \tau^r = \tau^{\ddagger} > \tau^{\dagger} = \delta = \lim_{\delta\to(\tau^{\dagger})^+} \tau^r$. ■

Proof of Proposition 4. First consider δ in the vicinity of zero. If $\delta = 0$, the firm chooses nondisclosure in equilibrium under both inspection policies (Proposition 1 and Proposition 2). Moreover, the equilibrium mean inspection intervals $\tau^p > 0$ and $\tau^r > 0$ with nondisclosure are both unique local minimizers of the corresponding long-run average social costs, denoted by $\overline{C}_p^N(\tau)$ and $\overline{C}_r^N(\tau)$ for periodic and random inspections, respectively. Under periodic inspections $\overline{C}_p^N(\tau) = \frac{h\lambda}{\mu} + \frac{\chi}{\tau} -$

$\frac{\lambda}{\mu}\left(\frac{\chi}{\tau}+h-r\right)\frac{1-e^{-\mu\tau}}{\mu\tau}$, obtained by applying the nondisclosure condition $s=0$ in (2). Under random inspections, $\overline{C}_r^N(\tau)=\frac{h\lambda}{\mu}+\frac{\chi}{\tau}-\frac{\lambda}{\mu}\left(\frac{\chi}{\tau}+h-r\right)\frac{1}{\mu\tau+1}$ (see (4)). From the inequality $1-e^{-x}>\frac{x}{x+1}$ we see that $\overline{C}_r^N(\tau)-\overline{C}_p^N(\tau)=\frac{\lambda}{\mu}\left(\frac{\chi}{\tau}+h-r\right)\left(\frac{1-e^{-\mu\tau}}{\mu\tau}-\frac{1}{\mu\tau+1}\right)>0$. Since $\overline{C}_p^N(\tau)<\overline{C}_r^N(\tau)$ for any given $\tau>0$ we have $\overline{C}_p^N(\tau^p)<\overline{C}_p^N(\tau^r)<\overline{C}_r^N(\tau^r)$, i.e., the minimum of $\overline{C}_p^N(\tau)$ is smaller than the minimum of $\overline{C}_r^N(\tau)$. Hence, $\overline{C}^p<\overline{C}^r$ when $\delta=0$. Now consider increasing δ infinitesimally from zero, which relaxes the maximum penalty constraint $\kappa\le r\delta$. Recall from Proposition 1 that this constraint binds in equilibrium under periodic inspections. Hence, increased δ leads to higher penalty which in turn enlarges the disclosure window s^p (see Proposition 1). From (2) we see that this change lowers the regulator's objective function $\overline{C}^p$. By contrast, under random inspections the constraint does not bind and $\overline{C}^r$ is independent of δ if δ is sufficiently small (see Proposition 2 and (5)). Hence, an infinitesimal increase in δ does not change $\overline{C}^r$. Combined with the earlier finding that $\overline{C}^p<\overline{C}^r$ when $\delta=0$, these observations imply that the difference $\overline{C}^r-\overline{C}^p$ becomes larger when δ is increased by a small amount. Therefore, $\overline{C}^p<\overline{C}^r$ for δ near zero. Next, consider δ away from zero. If δ is sufficiently large, the firm chooses full disclosure in equilibrium under both inspection policies (Proposition 1 and Proposition 2). Let $\overline{C}_p^F(\tau)$ and $\overline{C}_r^F(\tau)$ be the long-run average social costs under periodic and random inspections with full disclosure. Under periodic inspections $\overline{C}_p^F(\tau)=\frac{r\lambda}{\mu}-\frac{\chi\lambda}{2}+\frac{\chi}{\tau}\left(1-\frac{\lambda}{\mu}\right)$, obtained by applying the full disclosure condition $s=\tau$ in (2). Under random inspections, $\overline{C}_r^F(\tau)=\frac{r\lambda}{\mu}+\frac{\chi}{\tau}\left(1-\frac{\lambda}{\mu}\right)$ (see (4)). From Proposition 1 and Proposition 2 we see that the equilibrium mean inspection intervals with full disclosure are $\tau^p=\frac{1}{\mu}\ln(1+\mu\delta)$ and $\tau^r=\delta$. Thus, $\overline{C}^p=\overline{C}_p^F(\tau^p)=\frac{r\lambda}{\mu}-\frac{\chi\lambda}{2}+\frac{\chi\mu}{\ln(1+\mu\delta)}\left(1-\frac{\lambda}{\mu}\right)$ and $\overline{C}^r=\overline{C}_r^F(\tau^r)=\frac{r\lambda}{\mu}+\frac{\chi}{\delta}\left(1-\frac{\lambda}{\mu}\right)$ for given δ. Then we have $\overline{C}^p-\overline{C}^r=\chi\lambda\left(\frac{1}{\ln(1+\mu\delta)}-\frac{1}{\mu\delta}\right)\left[\left(\frac{\mu}{\lambda}-1\right)-\gamma(\mu\delta)\right]$, where $\gamma(y)\equiv\frac{y/2}{y/\ln(1+y)-1}$ is an increasing function. Recall the assumption $\mu\delta<\widehat{y}$ from §3.3, where $\widehat{y}$ uniquely solves the equation $\gamma(y)=\frac{\mu}{\lambda}-1$. Then $\gamma(\mu\delta)<\gamma(\widehat{y})=\frac{\mu}{\lambda}-1$; with $\mu\delta>\ln(1+\mu\delta)$, this implies $\overline{C}^p-\overline{C}^r>0$. Therefore, $\overline{C}^p>\overline{C}^r$ for sufficiently large δ. ■

Proof of Lemma 4. Let us use the subscripts p and r to denote periodic and random inspections, respectively. First consider the case where the firm chooses nondisclosure. This corresponds to $s^*=0$ under periodic inspections. Setting $s^*=0$ in the expressions in Lemma 1 yields $\overline{I}_p=\frac{\mu}{\mu\tau+\theta(\tau)}$ and $\overline{R}_p=\frac{\theta(\tau)}{\mu\tau+\theta(\tau)}$. Under random inspections, on the other hand, $\overline{I}_r=\frac{\mu}{\lambda+\mu}\left(\frac{1}{\tau}+\frac{\lambda}{\mu\tau+1}\right)$ and $\overline{R}_r=\frac{\lambda}{\lambda+\mu}\frac{1}{\mu\tau+1}$ (Lemma 3). In both cases, $\overline{C}=\chi\overline{I}+r\overline{R}+h\overline{B}=h\frac{\lambda}{\lambda+\mu}+\chi\overline{I}-(h-r)\overline{R}$ since $\overline{B}+\overline{R}=\frac{\lambda}{\lambda+\mu}$. From the inequalities $\frac{x}{x+1}<1-e^{-x}<x$, it follows that $\theta(t)$ defined in (1) satisfies $\frac{\lambda}{\lambda+\mu+1/t}<\theta(t)<\lambda t$. Then $\overline{I}_p-\overline{I}_r=\frac{\mu}{\mu\tau+\theta(\tau)}-\frac{\mu}{\lambda+\mu}\left(\frac{1}{\tau}+\frac{\lambda}{\mu\tau+1}\right)<\frac{\mu}{\lambda+\mu}\left(\frac{\lambda+\mu}{\mu\tau+\lambda(\lambda+\mu+1/\tau)}-\frac{\lambda+\mu+1/\tau}{\mu\tau+1}\right)=0$ and $\overline{R}_p-\overline{R}_r=\frac{1}{\mu\tau/\theta(\tau)+1}-\frac{\lambda}{\lambda+\mu}\frac{1}{\mu\tau+1}>\frac{\lambda}{\lambda+\mu}\left[\left(\mu\tau\left(\frac{1}{(\lambda+\mu)\tau}+1\right)+\frac{\lambda}{\lambda+\mu}\right)^{-1}-\frac{1}{\mu\tau+1}\right]=0$. Since $\overline{I}_p<\overline{I}_r$ and $\overline{R}_p>\overline{R}_r$, we have $\overline{C}_p<\overline{C}_r$, i.e., the long-run average social cost is lower under periodic inspections. Next, consider the case of full disclosure. This corresponds to $s^* = \tau$ under periodic inspections. Setting $s^* = \tau$ in the expressions in Lemma 1 yields $\overline{I}_p = \frac{\lambda\mu}{\lambda+\mu}\frac{e^{-\lambda\tau}}{1-e^{-\lambda\tau}}$ and $\overline{R}_p = \frac{\lambda}{\lambda+\mu}$. Under random inspections, on the other hand, $\overline{I}_r = \frac{\mu}{\lambda+\mu}\frac{1}{\tau}$, $\overline{R}_r = \frac{\lambda}{\lambda+\mu}$ (Lemma 3). Since $\frac{e^{-\lambda\tau}}{1-e^{-\lambda\tau}} < \frac{1}{\lambda\tau}$, we have $\overline{I}_p < \overline{I}_r$. Combined with $\overline{R}_p = \overline{R}_r$, this implies $\overline{C}_p < \overline{C}_r$. ■

B Performance Measures Under Periodic Inspections

In this section we derive the expectations of the performance measures X, I, R, B, and Ψ (summarized in Lemma 1) assuming that the firm employs the threshold disclosure policy in response to the regulator's periodic inspections. In the lemmas and proofs below, U and D refer to the exponential i.i.d. random variables with means $1/\lambda$ and $1/\mu$ representing the duration of each compliance and noncompliance episode, respectively. The Cases $1a$, $1b$, $2a$, $2b$, and $2c$ refer to the five possible outcomes in an inspection cycle, as described in §4.1.1 and in Figure 2.

B.1 Probability of Each Outcome in an Inspection Cycle

Lemma B.1 *The probability of each case in an inspection cycle is as follows: (i)* $\Pr(1a) = e^{-\lambda s}(1-\theta(\tau - s))$*; (ii)* $\Pr(1b) = (1-e^{-\lambda s})(1-\theta(\tau-s))$*; (iii)* $\Pr(2a) = e^{-\mu s}\theta(\tau - s)$*; (iv)* $\Pr(2b) = \frac{\mu}{\mu-\lambda}\left(e^{-\lambda s} - e^{-\mu s}\right)\theta(\tau - s)$*; (v)* $\Pr(2c) = \left(1 - \frac{\mu}{\mu-\lambda}e^{-\lambda s} + \frac{\lambda}{\mu-\lambda}e^{-\mu s}\right)\theta(\tau - s)$*, where* $\theta(t)$ *is defined in (1).*

Proof. In Cases $1a$ and $1b$ the firm is in compliance at both time zero and time $\tau - s$, an event that occurs with probability $1-\theta(\tau - s)$. On the other hand, in Cases $2a$, $2b$, and $2c$ the firm is in compliance at time zero but is in noncompliance at time $\tau - s$, an event that occurs with probability $\theta(\tau - s)$. Using the memoryless properties of these random variables, from Figure 2 we see that the probabilities for the five cases are: (i) $\Pr(1a) = \Pr(U > s)(1-\theta(\tau - s))$; (ii) $\Pr(1b) = \Pr(U \le s)(1-\theta(\tau - s))$; (iii) $\Pr(2a) = \Pr(D > s)\theta(\tau - s)$; (iv) $\Pr(2b) = \Pr(D \le s < D + U)\theta(\tau - s)$; (v) $\Pr(2c) = \Pr(D + U \le s)\theta(\tau - s)$. It remains to evaluate the conditional probabilities. Along with $\Pr(U > s) = e^{-\lambda s}$ and $\Pr(D > s) = e^{-\mu s}$, we have $\Pr(D \le s < D + U) = \frac{\mu}{\mu-\lambda}(e^{-\lambda s} - e^{-\mu s})$ and $\Pr(D + U \le s) = 1 - \frac{\mu}{\mu-\lambda}e^{-\lambda s} + \frac{\lambda}{\mu-\lambda}e^{-\mu s}$.
■

B.2 Expected Length of an Inspection Cycle

Lemma B.2 *The conditional expected length of an inspection cycle for each case in an inspection cycle is as follows: (i)* $E[X|1a] = \tau$*; (ii)* $E[X|1b] = \tau - \frac{s}{1-e^{-\lambda s}} + \frac{1}{\lambda} + \frac{1}{\mu}$*; (iii)* $E[X|2a] = \tau + \frac{1}{\mu}$*; (iv)* $E[X|2b] = \tau$*; (v)* $E[X|2c] = \tau - \frac{(\mu-\lambda)s - (e^{-\lambda s} - e^{-\mu s})}{\mu(1-e^{-\lambda s}) - \lambda(1-e^{-\mu s})} + \frac{1}{\lambda} + \frac{2}{\mu}$*. Unconditioning, the expected length of an inspection cycle is* $E[X] = \tau - s + \left(\frac{1}{\lambda} + \frac{1}{\mu}\right)(1-e^{-\lambda s}) + \left(\frac{1}{\mu} - \frac{\lambda+\mu}{\mu(\mu-\lambda)}(e^{-\lambda s} - e^{-\mu s})\right)\theta(\tau - s)$*.*

Proof. In Case $1a$ and Case $2b$ a cycle takes τ time units as no disclosure or detection occurs by then. In Case $2a$ the cycle that starts at time 0 ends after noncompliance is detected at time τ and until compliance is restored, which in total takes $\tau + 1/\mu$ time units in expectation. (The residual noncompliance duration after τ is exponentially distributed with mean $1/\mu$ due to memorylessness of D.) Consider Case 1b. The cycle lasts at least $\tau - s$ time units, after which it lasts until the noncompliance episode that starts between times $\tau - s$ and τ concludes. Let Z be the residual compliance duration starting from time $\tau - s$. Then Z is exponentially distributed with mean $1/\lambda$ because of memorylessness of U. The expected duration from time $\tau - s$ until compliance restoration conditional on $Z \le s$ is $E[Z|Z \le s] + E[D] = \int_0^\infty z f_Z(z|Z \le s)\,dz + \frac{1}{\mu} = \int_0^s z \frac{\lambda e^{-\lambda z}}{1-e^{-\lambda s}} dz + \frac{1}{\mu} = \frac{1}{\lambda} - \frac{s e^{-\lambda s}}{1-e^{-\lambda s}} + \frac{1}{\mu}$, where $f_Z(z|Z \le s) = \frac{f_Z(z)}{F_Z(s)}$ defined on the support $[0, s]$ is the truncated distribution of Z (Mood et al. 1973, p. 124). Hence, the expected cycle length of Case 1b is $\tau - s + E[Z|Z \le s] + E[D] = \tau - \frac{s}{1-e^{-\lambda s}} + \frac{1}{\lambda} + \frac{1}{\mu}$. Finally, consider Case 2c. The cycle lasts at least $\tau - s$ time units, after which it lasts until the state makes three transitions: to compliance and then to noncompliance between times $\tau - s$ and τ, and then back to compliance again. Let Y be the residual noncompliance duration starting from time $\tau - s$ which is exponentially distributed with mean $1/\mu$. Noting that the pdf and the cdf of the random variable $Y + U$ are $f_{Y+U}(z) = \frac{\lambda\mu}{\mu-\lambda}\left(e^{-\lambda z} - e^{-\mu z}\right)$ and $F_{Y+U}(z) = 1 - \frac{\mu e^{-\lambda z} - \lambda e^{-\mu z}}{\mu - \lambda}$ (Ebeling 2009, p. 235), the expectation of the truncated random variable $Y + U \le s$ is $E[Y + U|Y + U \le s] = \int_0^\infty z f_{Y+U}(z|Y + U \le s)\,dz = \int_0^s z \frac{f_{Y+U}(z)}{F_{Y+U}(s)} dz = \frac{e^{-\lambda s} - e^{-\mu s} - s(\mu e^{-\lambda s} - \lambda e^{-\mu s})}{\mu(1-e^{-\lambda s}) - \lambda(1-e^{-\mu s})} + \frac{1}{\lambda} + \frac{1}{\mu}$. Hence, the conditional expected length of a renewal in this case is $\tau - s + E[Y + U|Y + U \le s] + E[D] = \tau - \frac{(\mu-\lambda)s - (e^{-\lambda s} - e^{-\mu s})}{\mu(1-e^{-\lambda s}) - \lambda(1-e^{-\mu s})} + \frac{1}{\lambda} + \frac{2}{\mu}$. Finally, $E[X]$ is obtained by unconditioning the conditional expectations using the probabilities computed in Lemma B.1 and collecting terms. ■

B.3 Expected Number of Inspections Performed in an Inspection Cycle

Lemma B.3 *The expected number of inspections performed in an inspection cycle is* $E[I] = e^{-\lambda s} + \frac{\lambda}{\mu-\lambda}(e^{-\lambda s} - e^{-\mu s})\theta(\tau - s)$.

Proof. Since exactly one inspection is performed in Cases 1a, 2a, and 2b whereas none in Cases 1b and 2c (as a disclosure suspends inspections), the expected number of inspections performed per cycle is $\Pr(1a) + \Pr(2a) + \Pr(2b)$. Then the result follows from Lemma B.1. ■

B.4 Expected Duration of Suspended Noncompliance in an Inspection Cycle

Lemma B.4 *The expected duration of suspended noncompliance in an inspection cycle is* $E[R] = \frac{1-e^{-\lambda s}}{\mu} + \frac{\mu e^{-\mu s} - \lambda e^{-\lambda s}}{\mu(\mu-\lambda)}\theta(\tau - s)$.

Proof. Noncompliance is reported either by a detection or a disclosure. Thus, only Cases 1b, 2a, and 2c are relevant. Because of memorylessness of D, in each case the expected duration of a residual noncompliance duration after a report is equal to $1/\mu$. Accounting for all cases, $E[R] = \frac{1}{\mu}[\Pr(1b) + \Pr(2a) + \Pr(2c)]$. Then the result follows from Lemma B.1. ■

B.5 Expected Duration of Unsuspended Noncompliance in an Inspection Cycle

Lemma B.5 *The expected duration of unsuspended noncompliance in an inspection cycle is* $E[B] = \frac{\lambda}{\lambda+\mu}(\tau-s) + \left(\frac{1-e^{-\mu s}}{\mu} - \frac{1}{\lambda+\mu}\right)\theta(\tau-s)$.

Proof. The expected duration of noncompliance between time 0 and time $\tau - s$ is (Nakagawa 2005, p. 45) $E\left[\int_0^{\tau-s}\mathbf{1}(\text{noncompliance at } t)dt\right] = \int_0^{\tau-s}E[\mathbf{1}(\text{noncompliance at } t)]\,dt = \int_0^{\tau-s}\Pr(\text{noncompliance at } t)$ $= \int_0^{\tau-s}\theta(t)dt = \frac{\lambda}{\lambda+\mu}(\tau-s) - \frac{1}{\lambda+\mu}\theta(\tau-s)$. Since the noncompliance episodes that start before time $\tau - s$ is unreported, this expected duration is common in all five cases of Figure 2. Now consider the disclosure window $[\tau-s,\tau]$. In Cases $1a$, no noncompliance exists in the window, contributing zero to $E[B]$. In Case $1b$, any noncompliance episode that starts within the window is reported; hence, this case does not contribute to $E[B]$. In Case $2a$, noncompliance duration is greater than the window; hence, this case contributes s to $E[B]$. In Cases $2b$ and $2c$, the conditional expected duration of noncompliance in $[\tau-s,\tau]$ is equal to $E[D|D<s] = \int_0^s x\frac{\mu e^{-\mu x}}{1-e^{-\mu s}}dx = \frac{1}{\mu} - \frac{se^{-\mu s}}{1-e^{-\mu s}}$. Unconditioning using Lemma B.1 and adding the first result above yields $E[B] = \int_0^{\tau-s}\theta(t)dt + s\Pr(2a) + E[D|D<s](\Pr(2b)+\Pr(2c)) = \frac{\lambda}{\lambda+\mu}(\tau-s) + \left(\frac{1-e^{-\mu s}}{\mu} - \frac{1}{\lambda+\mu}\right)\theta(\tau-s)$. ■

B.6 Firm's Expected Cost in an Inspection Cycle

Lemma B.6 *The firm's conditional expected profit per inspection cycle for each case is as follows: (i)* $E[\Pi|1a] = r\tau$; *(ii)* $E[\Pi|1b] = r\left(\tau - \frac{s}{1-e^{-\lambda s}} + \frac{1}{\lambda}\right)$; *(iii)* $E[\Pi|2a] = r\tau - \kappa$; *(iv)* $E[\Pi|2b] = r\tau$; *(v)* $E[\Pi|2c] = r\left(\tau - \frac{(\mu-\lambda)s - \left(e^{-\lambda s} - e^{-\mu s}\right)}{\mu\left(1-e^{-\lambda s}\right) - \lambda\left(1-e^{-\mu s}\right)} + \frac{1}{\lambda} + \frac{1}{\mu}\right)$. *Unconditioning, the firm's expected cost per inspection cycle (including the penalty and the opportunity costs) is* $E[\Psi] = rE[X] - E[\Pi] = r\frac{1-e^{-\lambda s}}{\mu} + \left(r\frac{\mu e^{-\mu s} - \lambda e^{-\lambda s}}{\mu(\mu-\lambda)} + \kappa e^{-\mu s}\right)\theta(\tau-s)$.

Proof. In Case $1a$ and Case $2b$ production continues until time τ, when the inspection cycle ends with no disclosure or detection. Hence the expected profit is $r\tau$. In Case $1b$ production continues until the firm discloses noncompliance at its onset which then lasts D additional amount of time before the cycle ends. Therefore, production expects to last $E[X|1b] - \frac{1}{\mu}$ and consequently the expected profit is $r\left(E[X|1b] - \frac{1}{\mu}\right)$. Similarly, in Case $2c$ the expected profit is $r\left(E[X|2c] - \frac{1}{\mu}\right)$. In Case $2a$ production continues until time τ, when noncompliance is detected. Upon detection, the firm pays the fixed penalty κ. Therefore, the expected profit in Case $2a$ is $r\tau - \kappa$. The expressions for Cases $1b$ and $2c$ are obtained using the results in Lemma B.2. Finally, $E[\Psi] = rE[X] - E[\Pi]$ is obtained by unconditioning the conditional expectations using the probabilities computed in Lemma B.1 and collecting terms. ■

Acknowledgements

The author thanks Ed Kaplan, Serguei Netessine, Erica Plambeck, Robert Swinney, and Terry Taylor for helpful comments.

References

[1] Agrawal, V., M. Fergerson, L. B. Toktay, V. Thomas. 2012. Is leasing greener than selling? *Management Science*, 58(3), 523-533.

[2] Alizamir, S., F. de Véricourt, P. Sun. 2012. Efficient feed-in-tariff policies for renewable energy technologies. Working paper, Duke University.

[3] Avci, B., K. Girotra, S. Netessine. 2012. Electric vehicles with a batter switching station: Adoption and environmental impact. Working paper, INSEAD.

[4] Barlow R. E. and F. Proschan. 1996. Mathematical Theory of Reliability. SIAM, Philadelphia, PA

[5] Beavis, B., M. Walker. 1983. Random wastes, imperfect monitoring and environmental quality standards. *Journal of Public Economics*, 21, 377-387.

[6] Becker, G. S. 1968. Crime and punishment: An economic approach. *Journal of Political Economy*, 72(2), 169-217.

[7] Cachon, G. P. 2012. Supply chain design and the cost of greenhouse gas emissions. Working paper, University of Pennsylvania.

[8] Cachon, G. P., P. H. Zipkin. Competitive and cooperative inventory policies in a two-stage supply chain. *Management Science*, 45(7), 936-953.

[9] Caro, F., C. Corbett, T. Tan, R. Zuidwijk. 2012. Double-counting of emissions in carbon-optimal and carbon-neutral supply chains. Working paper, UCLA.

[10] Cohen, M. A. 1999. Monitoring and enforcement of environmental policy. *International Yearbook of Environmental and Resource Economics 1999/2000*, eds. Tietenberg, T., H. Folmer, 44-106. Edward Elgar Publishers.

[11] Drake, D., P. R. Kleindorfer, L. N. Van Wassenhove. 2012. Technology choice and capacity portfolios under emissions regulation. Working paper, Harvard Business School.

[12] Goodstein, E. S. 2011. *Economics and the Environment*. Wiley.

[13] Guo, R., H. Lee, R. Swinney. 2012. The impact of supply chain structure on responsible sourcing. Working paper, Stanford University.

[14] Heyman, D. P., M. J. Sobel. 1982. *Stochastic Models in Operations Research, Vol. 1: Stochastic Processes and Operating Characteristics*. Dover, Mineola, New York.

[15] Innes, R. 1999. Remediation and self-reporting in optimal law enforcement. *Journal of Public Economics*, 72(3), 379-393.

[16] Jira, C., M. W. Toffel. 2012. Engaging supply chains in climate change. *Manufacturing & Service Operations Management*, forthcoming.

[17] Kalkanci, B., E. Ang, E. Plambeck. 2012. Measurement and improvement of social and environmental performance under voluntary versus mandatory disclosure. Working paper, MIT.

[18] Kaplan, E. H. 2010. Terror queues. *Operations Research*, 58(4), 773-784.

[19] Kaplow, L., S. Shavell. 1994. Optimal law enforcement with self-reporting of behavior. *Journal of Political Economy*, 102(3), 583-606.

[20] Karlin, S., H. M. Taylor. 1975. *A First Course in Stochastic Processes.* Academic Press.

[21] Kim, S.-H., M. A. Cohen, S. Netessine, S. Veeraraghavan. 2010. Contracting for infrequent restoration and recovery of mission-critical systems. *Management Science*, 56(9), 1551-1567.

[22] Kim, S.-H., B. Tomlin. 2012. Guilt by association: Strategic failure prevention and recovery capacity investments. *Management Science*, forthcoming.

[23] Kolstad, C. D. 2011. *Environmental Economics*. Oxford University Press, New York.

[24] Kulkarni, V. G. 1995. *Modeling and Analysis of Stochastic Systems*. Chapman & Hall.

[25] Lanjouw, P. 2006. Small-scale industry, poverty, and the environment: A case study of Ecuador. Blackman, A., eds. *Small Firms and the Environment in Developing Countries: Collective Impacts, Collective Action.* RFF Press, Washington, DC, 47-71.

[26] Lee, D. 2011. Turning waste into by-product. *Manufacturing & Service Operations Management*, 13(1), 115-127.

[27] Lobel, R., G. Perakis. 2011. Consumer choice model for forecasting demand and designing incentives for solar technology. Working paper, University of Pennsylvania.

[28] Malik, A. S. 1993. Self-reporting and the design of policies for regulating stochastic pollution. *Journal of Environmental Economics and Management*, 24, 241-257.

[29] Nakagawa, T. 2005. *Maintenance Theory of Reliability.* Springer, London.

[30] Plambeck, E. L., T. A. Taylor. 2010. Competitors as whistle-blowers in enforcement of product standards. Working paper, Stanford University.

[31] Plambeck, E. L., T. A. Taylor. 2012. Supplier evasion of a buyer's audit: Implications for auditing and compliance with labor and environmental standards. Working paper, Stanford University.

[32] Rausand, M., M. A. Høyland. 2004. *System Reliability Theory: Models, Statistical Methods, and Applications*. Wiley.

[33] Ross, S. M. 1996. *Stochastic Processes*. Wiley.

[34] Russell, C. S. 1990. Game models for structuring monitoring and enforcement systems. *Natural Resources Modeling*, 4(2), 143-173.

[35] Short, J. L., M. W. Toffel. 2007. Coerced confessions: Self-policing in the shadow of the regulator. *Journal of Law, Economics, & Organization*, 24(1), 5-71.

[36] Su, S. Y. H., I. Koren, Y. K. Malaiya. 1978. A continuous-parameter Markov model and detection procedures for intermittent faults. *IEEE Transactions on Computers*, C-27(6), 567-570.

[37] Sunar, N., E. L. Plambeck. 2011. Greenhouse gas emissions accounting: Allocating emissions from process to co-products. Working paper, Stanford University.

[38] Tijms, H. C. 2003. *A First Course in Stochastic Models*. Wiley, UK.

[39] Toffel, M. W., J. L. Short. 2011. Coming clean and cleaning up: Does voluntary self-reporting indicate effective self-policing? *Journal of Law and Economics*, 54(3), 609-649.

[40] Wolff, R. W. 1989. *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, New Jersey.

[41] Zheng, Y.-S. 1992. On properties of stochastic inventory systems. *Management Science*, 38(1), 87-103.

Technical Appendix

Time to Come Clean? Disclosure and Inspection Policies for Green Production

C Additional Results and Proofs

Lemma C.1 *(For the proof of Proposition 1) Let* $\varphi(x) \equiv \frac{1}{x^3}\left((2-b)e^x - 2 - 2x - x^2\right)(3-x)$ *where* $0 \le b \le 2$. *Then* $\varphi(x) < 1$ *for all* $x > 0$.

Proof. Suppose $b = 0$. Applying l'Hopital's rule repeatedly,

$$\begin{aligned}
\lim_{x\to 0}\varphi(x) &= \lim_{x\to 0}\frac{2(e^x-1-x)(3-x)-2(e^x-1-x)+x^2}{3x^2} = \tfrac{1}{3} + \lim_{x\to 0}\frac{(4-2x)(e^x-1-x)}{3x^2} = \tfrac{1}{3} + \lim_{x\to 0}\frac{-2(e^x-1-x)+(4-2x)(e^x-1)}{6x} \\
&= \tfrac{2}{3} + \lim_{x\to 0}\frac{(2-2x)(e^x-1)}{6x} = \tfrac{2}{3} + \lim_{x\to 0}\frac{-2(e^x-1)+(2-2x)e^x}{6} = 1 - \tfrac{1}{3}\lim_{x\to 0} xe^x = 1
\end{aligned}$$

and similarly $\lim_{x\to\infty}\varphi(x) = 1 - \frac{1}{3}\lim_{x\to\infty} xe^x = -\infty$. Also, $\varphi'(x) = \frac{\zeta(x)}{x^4}$, where $\zeta(x) \equiv 18 + 8x + x^2 - \left(18 - 10x + 2x^2\right)e^x$. Observe

$$\begin{aligned}
\zeta(x) &= 18 + 8x + x^2 - \left(18 - 10x + 2x^2\right)\left(\textstyle\sum_{n=0}^{\infty}\frac{x^n}{n!}\right) \\
&= 18 + 8x + x^2 - 18\left(\textstyle\sum_{n=0}^{\infty}\frac{x^n}{n!}\right) + 10\left(\textstyle\sum_{n=0}^{\infty}\frac{x^{n+1}}{n!}\right) - 2\left(\textstyle\sum_{n=0}^{\infty}\frac{x^{n+2}}{n!}\right) \\
&= 18 + 8x + x^2 - 18\left(1 + x + \textstyle\sum_{n=0}^{\infty}\frac{x^{n+2}}{(n+2)!}\right) + 10\left(x + \textstyle\sum_{n=0}^{\infty}\frac{x^{n+2}}{(n+1)!}\right) - 2\left(\textstyle\sum_{n=0}^{\infty}\frac{x^{n+2}}{n!}\right) \\
&= x^2 + \textstyle\sum_{n=0}^{\infty}\left(\frac{-18}{(n+2)(n+1)} + \frac{10}{n+1} - 2\right)\frac{x^{n+2}}{n!} = x^2 - 2x^2\left(\textstyle\sum_{n=0}^{\infty}\frac{(n-1)^2}{(n+2)!}x^n\right) \\
&= x^2 - 2x^2\left(\tfrac{1}{2} + \textstyle\sum_{n=2}^{\infty}\frac{(n-1)^2}{(n+2)!}x^n\right) = -2x^2\textstyle\sum_{n=2}^{\infty}\frac{(n-1)^2}{(n+2)!}x^n < 0
\end{aligned}$$

for $x > 0$, implying $\varphi'(x) = \frac{\zeta(x)}{x^4} < 0$. Since $\varphi(x)$ starts from one at $x = 0$ and decreases to $-\infty$ as $x \to \infty$, we conclude $\varphi(x) < 1$ for all $x > 0$ if $b = 0$. Now suppose $0 < b \le 2$. Observe that the numerator of $\varphi(x)$ is equal to $-3b < 0$ at $x = 0$. Hence, $\lim_{x\to 0}\varphi(x) = -\infty$. Applying l'Hopital's rule repeatedly, we also have $\lim_{x\to\infty}\varphi(x) = 1 - \frac{2-b}{6}\lim_{x\to\infty} xe^x = -\infty$. Differentiating, $\varphi'(x) = \frac{1}{x^4}\left(18 + 8x + x^2 - (2-b)\left(9 - 5x + x^2\right)e^x\right)$. Therefore, $\varphi'(x) = 0$ is equivalent to $\beta(x) = 2 - b$ where $\beta(x) \equiv \frac{x^2+8x+18}{x^2-5x+9}e^{-x}$. Note $\beta(0) = 2$, $\lim_{x\to\infty}\beta(x) = 0$, and $\beta'(x) = -\frac{x^3(x+3)e^{-x}}{(x^2-5x+9)^2} < 0$. Since $\beta(x)$ decreases from 2 to zero as x goes from zero to infinity and $0 < b \le 2$, there is exactly one solution to $\beta(x) = 2 - b$. This implies that there is exactly one x that solves $\varphi'(x) = 0$, i.e., exactly one critical point of $\varphi(x)$ exists. Given that $\lim_{x\to 0}\varphi(x) = \lim_{x\to\infty}\varphi(x) = -\infty$, this critical point is a maximizer, which we denote as $\widehat{x}$. Evaluating $\varphi(x)$ at this point yields $\varphi(\widehat{x}) = \frac{1}{\widehat{x}^3}\left(\frac{\widehat{x}^2+8\widehat{x}+18}{\widehat{x}^2-5\widehat{x}+9} - 2 - 2\widehat{x} - \widehat{x}^2\right)(3-\widehat{x}) = \frac{(\widehat{x}-3)^2}{(\widehat{x}-3)^2+\widehat{x}} < 1$. Since the maximum of $\varphi(x)$ is smaller than one, we conclude $\varphi(x) < 1$ for all $x > 0$ if $0 < b \le 2$. ■

Lemma C.2 *(For §4.1.2) For fixed* τ, *if* $\kappa < \frac{r}{\mu}\left(e^{\mu\tau} - 1\right)$ *then* $\overline{I}^*$ *and* $\overline{B}^*$ *decrease in* κ *whereas* $\overline{R}^*$ *and* $\overline{\Psi}^*$ *increase in* κ. *If* $\kappa \ge \frac{r}{\mu}\left(e^{\mu\tau} - 1\right)$, *on the other hand,* $\overline{I}^*$, $\overline{R}^*$, $\overline{B}^*$, *and* $\overline{\Psi}^*$ *do not vary with* κ.

Proof. From Lemma 2 we have $s^* = \tau$ if $\tau \le \frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right)$ or equivalently $\kappa \ge \frac{r}{\mu}\left(e^{\mu\tau} - 1\right)$. Substituting this in the expressions in Corollary 1 yields $\overline{I}^* = \left(1 - \frac{\lambda}{\mu}\right)\frac{1}{\tau} - \frac{\lambda}{2}$, $\overline{R}^* = \frac{\lambda}{\mu}$, $\overline{B}^* = 0$, and $\overline{\Psi}^* = \frac{r\lambda}{\mu}$. Notice that all measures are independent of κ. If $\tau > \frac{1}{\mu}\ln\left(1 + \frac{\kappa\mu}{r}\right)$ or equivalently $\kappa < \frac{r}{\mu}\left(e^{\mu\tau} - 1\right)$, on the other hand, $s^* = \frac{1}{\mu} \ln\left(1 + \frac{\kappa\mu}{r}\right)$. Substituting this in the expressions in Corollary 1,

$$\begin{aligned}
\overline{I}^* &= \tfrac{1}{\tau} - \tfrac{\lambda}{\mu\tau} \ln\left(1 + \tfrac{\kappa\mu}{r}\right) - \tfrac{\lambda}{\mu^2\tau^2} \left(\ln\left(1 + \tfrac{\kappa\mu}{r}\right) - \tfrac{1}{2} \left(\ln\left(1 + \tfrac{\kappa\mu}{r}\right)\right)^2 + \tfrac{r}{r+\kappa\mu} - e^{-\mu\tau} \right), \\
\overline{R}^* &= \tfrac{\lambda}{\mu^2\tau} \left(\ln\left(1 + \tfrac{\kappa\mu}{r}\right) + \tfrac{r}{r+\kappa\mu} - e^{-\mu\tau} \right), \quad \overline{B}^* = \tfrac{\lambda}{\mu} - \tfrac{\lambda}{\mu^2\tau} \left(\ln\left(1 + \tfrac{\kappa\mu}{r}\right) + \tfrac{r}{r+\kappa\mu} - e^{-\mu\tau} \right), \\
\overline{\Psi}^* &= \tfrac{r\lambda}{\mu^2\tau} \left(1 + \ln\left(1 + \tfrac{\kappa\mu}{r}\right) - \tfrac{r+\kappa\mu}{r} e^{-\mu\tau} \right).
\end{aligned}$$

Differentiating $\overline{I}^*$ yields $\frac{\partial \overline{I}^*}{\partial\kappa} = \frac{\lambda}{\mu\tau^2} \frac{1}{r+\kappa\mu} \left(-\mu\tau - \frac{\kappa\mu}{r+\kappa\mu} + \ln\left(1 + \frac{\kappa\mu}{r}\right) \right) < -\frac{\lambda}{\mu\tau^2} \frac{\kappa\mu}{(r+\kappa\mu)^2} < 0$, where we used the condition $\tau > \frac{1}{\mu} \ln\left(1 + \frac{\kappa\mu}{r}\right)$ to establish the inequality. Differentiating $\overline{R}^*$ and $\overline{B}^*$ yields $\frac{\partial \overline{R}^*}{\partial\kappa} = -\frac{\partial \overline{B}^*}{\partial\kappa} = \frac{\lambda}{\mu\tau} \frac{\kappa\mu}{(r+\kappa\mu)^2} > 0$. Finally, differentiating $\overline{\Psi}^*$ yields $\frac{\partial \overline{\Psi}^*}{\partial\kappa} = \frac{r\lambda}{\mu^2\tau} \frac{\mu}{r+\kappa\mu} \left(1 - \frac{r+\kappa\mu}{r} e^{-\mu\tau}\right) > 0$, where we used the condition $\tau > \frac{1}{\mu} \ln\left(1 + \frac{\kappa\mu}{r}\right)$ to establish the inequality. ■

Lemma C.3 *(For Proposition 1 in §4.1.2) The equilibrium exists in the region $\tau \geq \sigma$ and satisfies $\kappa = r\delta$. Consequently, the firm sets $s^* = \sigma$ in equilibrium.*

Proof. Note that the social cost $\overline{C}^* = \chi\overline{I}^* + r\overline{R}^* + h\overline{B}^*$ can be written as $\overline{C}^* = \frac{h\lambda}{\mu} + \chi\overline{I}^* - (h-r)\,\overline{R}^*$ using the identity $\overline{B} = \frac{\lambda}{\mu} - \overline{R}$. Let $s(\kappa) \equiv \frac{1}{\mu} \ln\left(1 + \frac{\kappa\mu}{r}\right)$. We first prove that the equilibrium does not exist for $\tau < \sigma$. Suppose $\tau < \sigma$ and divide it into two regions: $\tau < s(\kappa) \leq \sigma$ and $s(\kappa) \leq \tau < \sigma$. The upper bound $s(\kappa) \leq \sigma$ is imposed by the maximum penalty constraint $\kappa \leq r\delta$. If $\tau < s(\kappa) \leq \sigma$, according to Lemma 2 the firm chooses $s^* = \tau$ and thus the social cost (2) reduces to $\overline{C}^* = \frac{\lambda}{\mu}\left(r - \frac{\chi\mu}{2}\right) + \left(1 - \frac{\lambda}{\mu}\right)\frac{\chi}{\tau}$, which is independent of κ but decreases in τ. Since $\overline{C}^*$ keeps decreasing in τ in the considered region $\tau < s(\kappa)$, the minimum of $\overline{C}^*$, if it exists in $\tau < \sigma$, should be found in the next region $s(\kappa) \leq \tau < \sigma$. In this region the firm chooses $s^* = s(\kappa)$ according to Lemma 2. In Lemma C.2 we proved that $\overline{I}^*$ with $s^* = s(\kappa)$ decreases in κ while $\overline{R}^*$ increases in κ. Consequently $\overline{C}^*$ decreases in κ if $s(\kappa) \leq \tau < \sigma$, and because $s(\kappa)$ increases, the minimum of $\overline{C}^*$, if it exists in this region, should satisfy $s(\kappa) = \tau$ where the maximum κ is found. But from Lemma 2 we see that $s(\kappa) = \tau$ implies $s^* = \tau$, in which case the social cost (2) again reduces to $\overline{C}^* = \frac{\lambda}{\mu}\left(r - \frac{\chi\mu}{2}\right) + \left(1 - \frac{\lambda}{\mu}\right)\frac{\chi}{\tau}$. Since this function decreases in τ for $\tau < \sigma$, by the similar argument as above the minimum does not exist in $\tau < \sigma$, thus confirming the statement we set out to prove. Now assume $\tau \geq \sigma$. With the maximum penalty constraint this condition requires $s(\kappa) \leq \sigma \leq \tau$. Since $s(\kappa) \leq \tau$, from Lemma 2 we see that the firm chooses $s^* = s(\kappa)$ and that $\overline{C}^*$ decreases in κ, by the same reasoning as above. Then the minimum of $\overline{C}^*$ satisfies $s(\kappa) = \sigma$ in the considered region $s(\kappa) \leq \sigma \leq \tau$ since $\overline{C}^*$ keeps decreasing as κ goes up until $s(\kappa)$, which increases, reaches its upper bound σ. Hence, the equilibrium exists for $\tau \geq \sigma$ and it satisfies $s(\kappa) = \sigma$ or equivalently $\kappa = r\delta$, at which $s^* = s(\kappa) = \sigma$. ■

Proposition C.1 *The firm employing the threshold disclosure policy under periodic inspections strictly prefers immediate disclosure to delayed disclosure. Similarly, the firm employing the binary disclosure policy under random inspections prefers immediate disclosure to delayed disclosure.*

Proof. (a) (Threshold disclosure policy under periodic inspections) Let *delayed threshold disclosure policy* be the modified version of the threshold disclosure policy under which the firm discloses noncompliance *after* its onset. Specifically, this policy works as follows: for any new noncompliance episode starting within the disclosure window, the firm discloses it either at $z > 0$ time units after its onset or at its conclusion, whichever happens first, unless it is first detected in the next scheduled inspection. If the inspection arrives before both of these two disclosure opportunities, then the firm does not disclose it until being detected. We prove the proposition in three steps. First, we identify the optimal policy among many possible delayed disclosure ("DD") policies with different delay times.

Second, we show that this optimal DD policy is weakly dominated by the nondisclosure ("ND") policy, i.e., the threshold policy of Definition 1 with $s = 0$. Finally, we show that ND policy is strictly dominated by the policy of Definition 1 with $s > 0$ ("original policy"). Without loss of generality we examine an arbitrarily chosen inspection cycle ("cycle"), whose start time is set to zero. Suppose that the disclosure window has a positive size, i.e., $s > 0$, and that a new noncompliance episode occurs within the window. (We only consider such an instance because, otherwise, there is no noncompliance to disclose under the threshold rule and therefore different disclosure policies lead to the same result.) Let $u \in (\tau - s, \tau)$ be the start time of the first such instance and D be its duration, which is exponentially distributed with mean $1/\mu$. Since D is random, the noncompliance end time $u + D$ may or may not be before time τ, when the next inspection is scheduled. Then under DD policy with the delay z, this noncompliance episode is reported to the regulator at time $u + \min\{D, z, \tau - u\}$; it is disclosed if $\min\{D, z\} < \tau - u$ whereas it is detected otherwise. Consider two cases: $z < \tau - u$ and $z \geq \tau - s$. Suppose $z < \tau - u$. Then $\min\{D, z\} < \tau - u$ and therefore the firm discloses the episode at time $u + \min\{D, z\}$. Since production is suspended immediately after disclosure and the firm pays the penalty κ for late disclosure, his expected per-cycle profit is equal to $rE[u + \min\{D, z\}] - \kappa$. Now suppose $z \geq \tau - u$. Then the firm discloses noncompliance if $D < \tau - u$ while he does not disclose and get detected if $D \geq \tau - u$. Hence, noncompliance is reported at time $u + \min\{D, \tau - u\}$. As before, production is suspended immediately after the report (via either disclosure or detection) and the firm pays the penalty; thus, the expected per-cycle profit is $rE[u + \min\{D, \tau - u\}] - \kappa$ in this case. Evaluating the expectations and combining the results, we find that, conditional on $u \in (\tau - s, \tau)$, the firm makes the expected per-cycle profit of $ru + \frac{r}{\mu}\left(1 - e^{-\mu \min\{z, \tau - u\}}\right) - \kappa$. Notice that this expression increases in z for $z < \tau - u$ until it becomes a constant afterwards, implying that under DD policy, the expected per-cycle profit conditional on $u \in (\tau - s, \tau)$ is maximized at any z satisfying $z \geq \tau - u$. We now show that setting $z \geq \tau - u$ also maximizes the firm's long-run average profit. Observe the following from Figure 2: (i) only Cases 1b and 2c in the figure are impacted by replacing the original policy with DD policy because the event we considered so far, i.e., a new noncompliance episode starting within the disclosure window, arises only in those two cases; (ii) the cycle length in each of the two cases remains unchanged because any new noncompliance episode starting within the disclosure window is eventually reported under DD policy, thus resetting the start of the next cycle to the end time of the episode just as under the original policy; (iii) the probability of each of the two cases remains unchanged because the underlying stochastic process is independent of the employed disclosure policy. Recall that, by the renewal-reward theorem, the firm's long-run average profit is equal to the ratio of the expected per-cycle profit and the expected cycle length; since the above observations imply that the denominator of this ratio is unchanged while the numerator is maximized by having $z \geq \tau - u$ for any realized value of $u \in (\tau - s, \tau)$, we see that the long-run average profit is also maximized when z satisfies the same condition. This condition states that the firm discloses the new noncompliance episode at or after the next scheduled inspection ($u + z \geq \tau$) regardless of its start time u, unless it concludes before then ($u + D < \tau$). Note that the firm is indifferent between detection and disclosure at time τ because in each case he incurs the penalty κ and earns the identical revenue by time τ. Adopting the tie-breaking convention of detection over disclosure in such a case, we can summarize the optimal policy within the class of DD policies as follows: if a new noncompliance episode starting within the disclosure window concludes before the next scheduled inspection, then disclose it at its conclusion; otherwise, do not disclose it and wait until being detected in the next inspection. Next, we prove that the policy we just stated is weakly dominated by ND policy. As before, the only cases in Figure 2 that are impacted by the choice of a disclosure policy are Cases 1b and 2c. For these cases, under the optimal DD policy described above, noncompliance is detected at time τ if $D \geq \tau - u$ for any given $u \in (\tau - s, \tau)$, whereas it is disclosed at time $u + D$ if $D < \tau - u$. Under ND policy, on the other hand, nondisclosure is detected at time τ if $D \geq \tau - u$ whereas it is never disclosed if $D < \tau - u$. Therefore, the two policies differ only by whether or not noncompliance is disclosed when $D < \tau - u$:

under the optimal DD policy the firm discloses noncompliance at its conclusion and pay the penalty, whereas under ND policy the firm does not disclose it and avoids the penalty. The corresponding per-cycle profits are $r(u+D)-\kappa$ and $r\tau$, or equivalently the corresponding per-cycle opportunity costs are $\kappa+r(\tau-u-D)>0$ and zero, respectively. In addition, the corresponding cycle lengths are $u+D$ and τ, respectively. Since the opportunity cost per unit time for Cases 1b and 2c is lower under ND policy than under the optimal DD policy (zero vs. positive) for any realized value of $u\in(\tau-s,\tau)$, the long-run average cost for the firm is also lower under ND policy. Hence, the firm prefers ND policy to the optimal DD policy and thus to any DD policy. Note that we arrived at this conclusion under the assumption $s>0$; if $s=0$, then there is no distinction between the policies because disclosure never happens. Therefore, ND policy weakly dominates DD policy. Finally, we show that ND policy is dominated by the original policy. Note that ND policy is identical to the original policy with s set to zero. In other words, the firm's long-run average cost under ND policy is exactly the same as $\overline{\Psi}$ appearing in Corollary 1 with $s=0$. But we proved in Lemma 2 that $\overline{\Psi}$ with $\kappa>0$ is minimized at $s>0$; hence, the policy that requires $s=0$ for $\overline{\Psi}$ (ND policy) is strictly dominated by the policy that does not (original policy). Then, since ND policy weakly dominates DD policy, the original policy strictly dominates DD policy.

(b) (Binary disclosure policy under random inspections) The unit of the analysis in this case is the compliance-noncompliance cycle of length $U+D$ (see the proof of Lemma 3). Under DD, the firm discloses noncompliance $z\in(0,D]$ time units after the onset, incurring the late disclosure penalty κ. Consider two cases: $T_r\le D$ and $T_r>D$. If $T_r\le D$, noncompliance is detected at time $U+T_r$ if $T_r\le z$ whereas it is disclosed at time $U+z$ if $T_r>z$. In each case the penalty κ is incurred and production is suspended as soon as noncompliance is reported. Hence, the conditional expected per-cycle profit is $rE[U]+rE[\min\{T_r,z\}]-\kappa$. If $T_r>D$, on the other hand, no detection occurs but noncompliance is disclosed at time $U+z$. The corresponding conditional expected per-cycle profit is $rE[U]+rz-\kappa$. Combining the two cases, the firm's expected profit per cycle is $rE[U]-\kappa+rE[\min\{T_r,z\}]\Pr(T_r\le D)+rz\Pr(T_r>D)$. Notice that this quantity increases in z. This implies that the firm maximizes its expected per-cycle profit by delaying disclosure as much as possible, i.e., disclose a noncompliance episode at its conclusion unless it is detected. However, this policy is dominated by ND since the latter results in a saving of the penalty κ, which is avoided if the firm does not disclose noncompliance at its conclusion, without affecting the revenue. Therefore DD is never optimal; it is dominated by ND. Note that, if $\kappa\ge r\tau$, DD is also dominated by FD since in this case FD dominates ND, which in turn dominates DD. ■

References

[1] Becker, G. S. 1968. Crime and punishment: An economic approach. *Journal of Political Economy*, 72(2), 169-217.

[2] Ebeling, C. E. 2009. *An Introduction to Reliability and Maintainability Engineering.* Waveland Press, IL.

[3] Kaplow, L., S. Shavell. 1994. Optimal law enforcement with self-reporting of behavior. *Journal of Political Economy*, 102(3), 583-606.

[4] Karlin, S., H. M. Taylor. 1975. *A First Course in Stochastic Processes.* Academic Press.

[5] Mood, A. M., F. A. Graybill, D. C. Boes. 1974. *Introduction to the Theory of Statistics.* McGraw-Hill.

[6] Nakagawa, T. 2005. *Maintenance Theory of Reliability.* Springer, London.