Journal of Public Economics 71 (1999) 415–440

Truth or consequences
Enforcing pollution standards with self-reporting

John Livernois*, C.J. McKenna

Department of Economics, University of Guelph, Guelph, Ontario, N1G 2W1, Canada

Received 1 July 1996; received in revised form 1 May 1997; accepted 2 June 1998

Abstract

Evidence suggests that a surprisingly large fraction of firms comply with pollution emission standards even though expected penalties for noncompliance are low. We offer an explanation of this puzzle by extending the standard model of enforcement to include a self-reporting requirement and enforcement power. These extensions are enough to challenge the conventional result that higher fines lead to higher compliance rates. We find that under plausible conditions, higher compliance rates are achieved with lower fines for noncompliance and the cost of enforcing a given level of aggregate pollution is minimized by setting the fine for noncompliance equal to zero. © 1999 Elsevier Science S.A. All rights reserved.

Keywords: Pollution control; Environmental regulation; Compliance; Self-reporting

JEL classification: D62; D82; L51

1. Introduction

Theory suggests that firms comply with pollution laws when the cost of compliance is less than the expected penalty for noncompliance and not otherwise. The empirical implication is that we should observe, ceteris paribus, low compliance rates when expected penalties are low and high compliance rates when

*Corresponding author. Tel.: +1 519 8244120, ext 8946; fax: +1 519 7638497; e-mail: live@css.uoguelph.ca

PII: S0047-2727(98)00082-6

expected penalties are high. It is puzzling therefore, that we commonly observe high compliance rates even when expected penalties are very low. For example, Russell et al. (1986) observe that expected penalties for noncompliance with pollution standards are low in the U.S. and yet, compliance rates are surprisingly high. Cropper and Oates (1992) concur that substantial compliance seems to exist even though few polluters have actually been fined for their violations. Moreover, they note, where such fines have been levied, they have been quite small.[1] Evidence from Canada shows a similar pattern. Data on industrial compliance rates over the period 1986–89 show that six of the nine industrial sectors were in compliance with all monthly pollution standards more than 70% of the time.[2] Yet plants that violate monthly pollution limits are rarely fined.[3]

If the expectation of being fined is low, why do firms comply with pollution standards? The purpose of this paper is to address this question. We construct a model that extends the standard enforcement model in two directions. The first is a requirement that firms monitor their own pollution emissions and report their compliance status to the enforcement agency. The second is to allow the enforcement agency the power to take enforcement actions.

Self-reporting is becoming an increasingly common feature of enforcement, particularly enforcement of pollution standards. For example, Title V of the U.S. Clean Air Act requires all major emission sources to evaluate their own compliance status and submit certified reports of compliance and noncompliance.[4] Enforcement power is also a common and, indeed, essential feature of environmental regulations. Regulators, such as the Environmental Protection Agency in the U.S., typically have a broad spectrum of tools to use in bringing enforcement actions against violators to return them to a state of compliance.[5]

Our paper is closest in purpose to Harrington (1988) who attempts to explain the puzzling phenomenon of high compliance rates and low expected fines by modelling enforcement as a repeated game between the agency and the firm. Just

[1] Cropper and Oates (1992) pages 696–7.

[2] Of the remaining three sectors, only one shows compliance rates below 50%. These data apply to the Province of Ontario and are taken from the Ontario Ministry of the Environment, 1986-89, Annual Direct Discharge Reports.

[3] No comprehensive data set exists on enforcement activity. However, available evidence taken from Statistics Canada, ''Human Activity and the Environment'', Catalogue 11-509E Occasional, indicates that convictions and fines are rare. For example, from April to September 1991, of the 1872 inspections conducted by Environment Canada, 49 warnings were issued to firms and yet only 6 were prosecuted. Over the 30 month period from April 1988 to September 1990, a total of 33 fines were levied in all of Canada by Environment Canada.

[4] Similar obligations exist in Canada. For example, the Environmental Protection Act of Ontario requires self-monitoring and self-reporting on a regular basis in a number of industrial sectors. In addition, the U.S. Environmental Protection Agency recently formalized the policy of offering reduced penalties to firms that voluntarily report a violation, even though they are not required to do so by law.

[5] Enforcement tools include warning letters, administrative control or compliance orders, civil complaints, and criminal indictments.

as insurance companies use experience-rating to better judge the risk class of their customers, Harrington allows the agency to use firms' past compliance behaviour to divide them into low- and high-risk (of violation) groups. Harrington shows that the compliance rate of firms in the second group is maximized by setting fines for noncompliance as high as possible for the second group but equal to zero (the lowest possible) for firms in the first group. As a result, his model predicts we should never observe first-time offenders being fined.

Our paper is closest in structure to Kaplow and Shavell (1994) and Malik (1993) who add a self-reporting requirement to the standard model of enforcement. In this setting, firms would confess to (self-report) violating a pollution standard only if they expected that not confessing would lead to a more severe punishment. If we let F_2 be the fine imposed on violators apprehended through random inspections that occur with probability π and F_1 be the fine imposed with certainty on firms that confess to a violation, we clearly need $F_1 \leq \pi F_2$ to induce firms to confess truthfully. Malik and Kaplow and Shavell show that the addition of a self-reporting requirement to the standard probabilistic model of enforcement unambiguously raises social welfare. One way to see this is that it lowers the cost of enforcement because it reduces the number of firms that need to be inspected.[6]

We offer a different explanation than Harrington (1988) of the puzzling observation of high compliance rates and low expected fines in this paper. While not denying the important leverage that memory gives the enforcement agency by allowing it to use experience rating, we abstract from it in this paper in order to focus on the different but, we believe, equally important enforcement tools of the self-reporting requirement and enforcement power.

Our model differs in structure from Malik (1993); Kaplow and Shavell (1994) because we allow for repeated interaction between regulator and firms and because we assume the regulator exploits the information revealed when firms self-report. In particular, we assume the regulator has the power to force any firm revealed to be in violation of a pollution law to return to compliance. This has the important implication that truth-telling is not incentive-compatible for all firms.

We find that the addition of self-reporting and enforcement power to the standard enforcement model is enough to challenge the conventional result that higher fines lead to higher compliance rates. Specifically, we show that under plausible conditions, compliance rates are a decreasing function of the fine for noncompliance and, that under these conditions, it is optimal to set the fine for noncompliance equal to zero, where the objective is to minimize the social cost of achieving a given pollution target. The implication is that observing low, even zero, fines for noncompliance in some polluting sectors, even though compliance

[6] In the context of pollution control, Swierzbinski (1994); Harford (1987) include self-reporting as a feature of their analysis; however, neither of these papers relate directly to our paper. In the context of tax compliance, Alm et al. (1992) find experimental evidence suggesting people truthfully report taxable income because they overweight the probability of inspection/audit.

is imperfect, is consistent with rational behaviour by firms and the enforcement agency.

In the next section, we introduce a model of pollution that allows for both deliberate violations of the pollution standard and violations which occur through random equipment failure. The third section introduces the self-reporting requirement and investigates the relationship between the aggregate compliance rate and the fine for noncompliance. In Section 4 we consider the consequences of allowing the enforcement agency to set the noncompliance fine so as to minimize the cost of enforcement. Extensions to the model are considered in Section 5 and conclusions are drawn in Section 6.

2. The pollution model

Firms are assumed to emit a single type of pollution and are required by law to respect an exogenous upper limit on allowable emissions, called the pollution standard. The responsibility of the enforcement agency (which we shall also call the regulator) is to enforce the pollution standard for individual firms; however, pollution emissions are unobservable to the regulator except during costly inspection visits when they can be accurately measured. Firms found in violation of the standard can be charged a fine. We assume there are no costs associated with levying fines.

We distinguish between initial and continuing compliance. A firm is said to be in initial compliance if it has installed a production process that makes compliance technically feasible. For concreteness, we assume the firm is in initial compliance if it has installed a pollution abatement device (PAD) and not otherwise. Without a PAD, compliance is feasible only by not operating the plant.

It is a relatively simple matter for the regulator to verify a firm's initial compliance status. Inspectors need only determine whether or not a firm's production process is technically capable of meeting the pollution standard, just as an automobile can be inspected to verify the presence of a catalytic converter. A firm that has installed a PAD is said to be in initial compliance.

A firm is said to be in continuing compliance if it actually operates the PAD and repairs it when it fails. It is more difficult for the regulator to verify a firm's continuing compliance status: while a single inspection visit can verify a firm's initial compliance status, repeated, nearly continuous, inspection visits would be required to determine with high probability a firm's continuing compliance status. We assume that all firms are in initial compliance and focus on the more challenging and interesting problem of continuing compliance in this paper.

The cost of operating the PAD is c per period. A firm's manager can bypass the PAD in production and thereby wilfully violate the pollution standard at a saving of c per period. However, violations of pollution standards probably arise more frequently from a lack of adequate maintenance and repairs to the PAD than from a deliberate act of violation. To capture this, we assume the PAD breaks down with probability β per period. To remain in compliance, the firm must immediately undertake repairs at a cost k which, we assume, leads to instantaneous resumption of its compliance status. If the PAD is not repaired, the firm's pollution emissions increase causing it to be in violation of the pollution standard. It continues to be in violation until the equipment is repaired.

Firms differ according to their compliance cost, c, but we assume that the repair cost, k, is the same for all firms. The distribution function of c, $g(c)$ defined over the interval $[0, \bar{c})$, is continuous, and is common information to firms and the regulator but the regulator does not know the compliance cost of any individual firm. We define the cumulative distribution function as $G(c)$.

At the beginning of each period, the firm discovers the state of the world that prevails. With probability $(1-\beta)$, its PAD is functioning and remains functioning for the duration of the period; with probability β, it is broken. If it is broken, it can be repaired immediately at cost k if the firm wishes to remain in compliance. It then returns to the working state for the duration of the period. If it is not repaired, it remains broken and the firm enters and remains in a state of non-compliance until it is repaired.

Our model contains, as a special case, the standard model of probabilistic enforcement if we rule out the possibility of equipment failure (set $\beta = 0$). Then, firms comply with the pollution standard if they operate the PAD at a cost c. Assuming risk neutrality, a firm complies if

$$c \leq \pi F \tag{1}$$

and does not comply otherwise, where π is the probability of inspection and F is the fine imposed on violators.[7] On average, the industry compliance rate is $G(\pi F)$.

In the context of pollution control, the enforcement models of Harford (1978); Storey and McCabe (1980) have this basic flavour although the probability of detection and the size of the fine are allowed to depend on the size of the violation in Harford's model. The industry compliance rate is increased by increasing either the fine, F, or the inspection probability, π. Optimal enforcement, as Becker (1968) first concluded, requires setting the fine at its maximum level, $\bar{F}$.[8] As noted earlier, this model is inconsistent with the observation of high compliance rates and low expected fines.

[7]This model of probabilistic enforcement is best thought of as a one-shot game. If the game were repeated over time, it would be necessary to assume the regulator can commit to not use the information it acquires when a firm is identified as a violator. For example, it must commit to not force the firm to return to compliance through higher sanctions or more intensive monitoring.

[8]The maximum level for the fine is assumed to be beyond the control of the enforcement agency.

3. Self-reporting

We assume that firms are required to monitor their own pollution emissions and submit compliance reports to the enforcement agency.[9] To verify the truthfulness of reports, the agency randomly inspects firms reporting compliance (there is no need to inspect firms reporting non-compliance). The penalty for submitting a false report is a fine, F_2 and an order to return to compliance. The penalty for truthfully reported noncompliance is a fine F_1 and an order to return to compliance.

The order to return to compliance is the only action we assume the regulator takes in response to the information acquired about a firm when it is identified as a type that does not comply.[10] Firms that have been ordered to return to compliance are allowed one period to comply with that order. We assume all firms delay their return to compliance to the full extent possible but that all do so by the end of the period. Although this assumption is a significant simplification of the enforcement process in practice,[11] we believe it captures an important element of enforcement not currently present in the literature.[12]

We define an *enforcement policy* as a triple $\{F_1, F_2, \pi\}$, $0 \leq F_i \leq \bar{F}$, $i = 1, 2$ and $\pi \in (0, 1)$. We initially study options for firms given an arbitrary enforcement policy and then look at the choice of policy facing the enforcement agency.

3.1. Strategy choices

For now, we assume the firm cannot choose to violate the pollution standard by simply not operating the PAD. Instead, violations can occur only after equipment failure. We relax this assumption in Section 5.

A firm can follow one of three strategies, S_{ij}, where $i = C$ or N for compliance

[9] Private monitoring costs are assumed to be zero as in Kaplow and Shavell (1994); Malik (1993).

[10] To keep the model tractable, we assume the regulator does not alter its policy of random inspections in response to information acquired about firms' types. The regulator has no incentive to alter its monitoring of firms that have violated the standard as a result of equipment failure and confess; but it does have an incentive to monitor more intensively firms that have been identified via inspections as violators, just as it does in the standard model of probabilistic enforcement. However, if we allowed for this, our model would mix features of Harrington's (Harrington, 1988) experience-rating model with our self-reporting model. As Harrington has already shown, this would bias our model in favour of providing an explanation for low fines. To avoid this bias, to keep the model tractable, and to focus on the power of self-reporting alone as a possible explanation of low fines, we assume the regulator cannot use experience-rating to adjust inspection rates.

[11] In practice, the enforcement agency is likely to issue a warning letter or a legally enforceable administrative order to the firm to return to compliance within a specified period of time. If the firm does not comply with this order, the agency may then proceed to take legal actions to restore compliance.

[12] The notable exception is the model of ''voluntary compliance'' by Russell et al. (1986). In practice, the order itself has to be enforced and this is done perhaps with the threat of higher and certain penalties. We take a simplified view of this control process which, we believe, nevertheless captures the essence of enforcement power.

or noncompliance and $j = T$ or F for truthful or false reporting of the state of compliance.

1. S_{CT}: full compliance *and* truth-telling means that the firm operates the PAD at a cost of $c > 0$ per period and repairs the machine (instantaneously) if it breaks, at a cost of $k > 0$. It is clear that there is no incentive for the fully-complying firm to be other than truthful, so that S_{CF} is ruled out.
2. S_{NT}: non-compliance and truth-telling means that the firm operates the PAD but does not repair it if it breaks down. Under S_{NT}, the firm truthfully reports that it is in a state of non-compliance and can expect certain punitive action.
3. S_{NF}: non-compliance and false reporting means that the firm operates the PAD, does not repair it if it breaks down, but always reports that it is in compliance. In this case, imperfect monitoring by the regulator means that the firm faces some chance of punitive action.

A firm will choose the strategy with the highest expected payoff, which clearly depends on the structure of penalties, costs of operation and repair of the PAD, the probabilities of failure and detection, and the magnitude of any fines imposed by the regulator.

A firm that reports a violation is assessed a fine of F_1 and is allowed to operate out of compliance for one period. At the end of the period, it repairs the PAD as required, at cost k. It then starts the next period, once again facing a probability β that the PAD is broken and probability $(1 - \beta)$ that it is not broken.[13]

We assume the regulator inspects firms randomly with the probability of a firm being inspected being $\pi \in (0, 1)$. The regulator is able to determine without error the compliance status of the firm at the time of inspection. Effectively then, an inspection determines whether a firm is telling the truth or not.[14] Clearly, the regulator does not inspect firms that report a violation, but only those firms claiming to be in compliance.

3.2. The optimal strategies

Let Z_{ij} be the expected discounted cost of following strategy S_{ij}. Although we set up the model as being dynamic, the horizon is infinite and all the stochastic processes are stationary so that a firm following a particular strategy at date t will

[13] Note that the PAD only occupies two states: functioning and broken so we are not concerned with issues of 'partial' compliance.

[14] In reality, the inspection determines whether or not the firm was falsifying records of monitoring and reporting. This might involve tips from disgruntled employees or direct evidence of falsification. However, we take a stylized view of this process by assuming the truth-telling is verified on inspection.

find it optimal to follow the same strategy at date $t+1$. The date subscript is therefore omitted from the calculations below.

The expected cost from adopting strategy S_{CT}, (full compliance), is given by

$$Z_{CT} = (1-\beta)c + \beta(c+k) + \delta Z_{CT}$$

With probability $1-\beta$, the PAD functions this period and the firm pays the cost of operating the PAD. With probability β the PAD breaks down, is repaired immediately at cost k and then is operated for the duration of the period at cost c. Next period, the firm expects to pay Z_{CT} again discounted one period by the discount factor δ. Note that though this firm reports compliance and may therefore be inspected, there are no cost-relevant implications of such an inspection, as long as the inspection does not interfere with the firm's operation and is totally accurate. Solving this gives

$$Z_{CT} = \frac{\beta k + c}{1-\delta} \tag{2}$$

The expected cost of adopting strategy S_{NT}, (non-compliance and truth-telling) is

$$Z_{NT} = (1-\beta)c + \beta(F_1 + \delta k) + \delta Z_{NT}$$

If the PAD does not fail, (probability $1-\beta$), cost is c as before. If the PAD breaks down, (probability β), a non-compliance report is filed, the firm is fined F_1 and is ordered to repair the PAD by the end of the period (hence the discounting factor on k). Note that the cost of operating the PAD, c, is avoided for one period. Next period, the process begins again yielding a present value of cost of δZ_{NT}. Solving gives

$$Z_{NT} = \frac{(1-\beta)c + \beta(F_1 + \delta k)}{1-\delta}$$

The expected cost of adopting strategy S_{NF}, (non-compliance and non-truth-telling), is more complex because inspection is both possible and will have payoff-relevant consequences. We have

$$Z_{NF} = (1-\beta)[c + \delta Z_{NF}] + \beta(\pi[(F_2 + \delta k) + \delta Z_{NF}] + (1-\pi)\delta Z_0 \tag{3}$$

where

$$Z_0 = \pi[(F_2 + \delta k) + \delta Z_{NF}] + (1-\pi)\delta Z_0 \tag{4}$$

If the machine does not break (with probability $1-\beta$), the firm complies and continues next period with the same expected cost discounted by one period. Note again that in this case, the firm will report compliance and, if inspected, will incur no additional costs. If the machine breaks down (with probability β), the firm does not report the violation and does not repair the machine. With probability π, the firm is inspected, fined F_2, forced to repair the machine next period and continues next period with the expected cost Z_{NF} discounted one period. If the firm is not inspected (with probability $1-\pi$) it incurs no cost of operating the broken PAD. Next period, the firm pays Z_0, which is a recursive formula repeating the same potential sequence of events of being inspected or not. It is the expected present-valued cost of starting the period with a broken PAD.

Solving for Z_{NF} gives

$$Z_{NF} = \left[\frac{(1-\beta)[1-\delta(1-\pi)]c + \beta\pi(F_2 + \delta k)}{(1-\alpha)(1-\delta)} \right] \tag{5}$$

where $\alpha \equiv \delta(1-\pi)(1-\beta) \in (0, 1)$.

Firms differ in their costs of operating the PAD and are distributed according to the continuous distribution function G, and density g, defined on $c \in [0, \bar{c}]$. A firm with a given c chooses the strategy with the lowest expected cost:

$$\min_S \{Z_{NT}, Z_{CT}, Z_{NF}\}$$

To determine how firms will be sorted among these three strategies in equilibrium, we need to identify the shape of the lower envelope of the cost functions for these three strategy choices. Inspection reveals that all cost functions, Z_{ij}, are linear and increasing in c.

In addition, it can be shown that

$$\frac{\partial Z_{CT}}{\partial c} > \frac{\partial Z_{NT}}{\partial c} > \frac{\partial Z_{NF}}{\partial c}$$

Given these properties, a number of possibilities arise for the shape of the lower envelope of the cost functions. However, the main features of the model are best illustrated by analyzing the case shown in Fig. 1. Here, we assume that each payoff line intersects the other two in the interior of $[0, \bar{c}]$.

Fig. 1 is drawn so that no strategy is dominated by another over all possible values of compliance cost, c. In this case, firms are sorted into three groups according to which of the three strategies is optimal over some range. For firms with low compliance costs, $0 < c < \hat{c}$, the optimal strategy is S_{CT}, full compliance. Firms with intermediate compliance costs, $\hat{c} < c < \tilde{c}$, choose to violate but truthfully report violations to the regulatory agency. Firms with high compliance costs, $\tilde{c} < c < \bar{c}$, choose to violate but report instead that they are in compliance.

Another situation is represented in Fig. 2, in which strategy S_{NT} is dominated by the other two strategies, with the critical value c' determining whether a firm finds it cost-effective to adopt strategy S_{CT} or strategy S_{NF}. Clearly the configurations in Figs. 1 and 2 are two of several possibilities which can arise depending on the parameter values of the model. However they are two cases particularly relevant for our purposes, and we return to them presently.

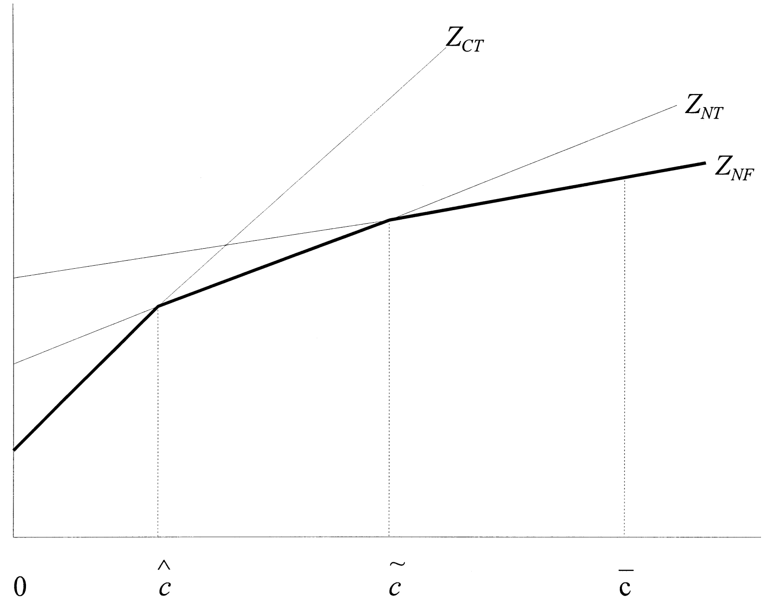

Fig. 1. The cost envelope with three types of firm.

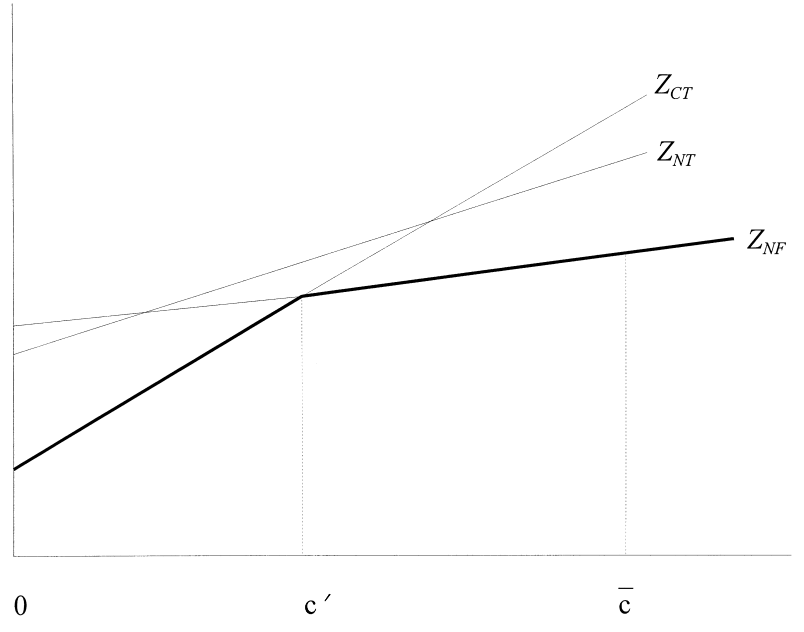

Fig. 2. The cost envelope with two types of firm.

Figs. 1 and 2 show that three critical values for c are going to play a role in our analysis. It is also apparent that these critical values will depend on the parameters of the enforcement policy as well as other parameters of the model. To find the value $\hat{c}$ in terms of the parameters of the model we clearly set $Z_{\mathrm{CT}} = Z_{\mathrm{NT}}$ and solve for c to give

$$\hat{c} = F_1 - k(1 - \delta) \tag{6}$$

It is obvious that $\hat{c}$ is increasing in F_1 and is independent of F_2 and π.[15] To find $\tilde{c}$, we set $Z_{\mathrm{NT}} = Z_{\mathrm{NF}}$ to give

$$\tilde{c} = \frac{\pi F_2 - (1 - \alpha)F_1 - \delta k(1 - \alpha - \pi)}{\alpha} \tag{7}$$

so that $\tilde{c}$ is decreasing in F_1, increasing in F_2 and, for $F_2 > F_1$, is increasing in π.[16]

Finally, we find c' by setting $Z_{\mathrm{CT}} = Z_{\mathrm{NF}}$:

$$c' = \pi(F_2 + \delta k) + k[1 - \delta(1 - \beta)(1 - \pi)] \tag{8}$$

so that c' is independent of F_1, increasing in F_2 and increasing in π.[17]

3.3. Truth-telling

In our model, firms choose not only compliance versus noncompliance, they also choose whether to file a truthful or false report. Our interest here is in determining for which firms truth-telling is incentive compatible. It is straightforward that firms that comply always tell the truth. However, firms that do not comply adopt the truth-telling strategy, S_{NT}, rather than the false-reporting strategy, S_{NF}, if and only if

$$Z_{\mathrm{NT}} \leq Z_{\mathrm{NF}} \tag{9}$$

which reduces to

$$F_1 + \delta k \leq \frac{\pi}{1 - \alpha}(F_2 + \delta k) - \frac{\alpha}{1 - \alpha}c$$

The truth-telling constraint in Eq. (9) is not necessarily satisfied for all firms. Even if the incentive for truth-telling is maximized by setting F_1 at its minimum level ($F_1 = 0$) and F_2 is at its maximum level ($F_2 = \bar{F}$), there will still exist some high values of $c \in [0, \bar{c}]$ which violate the constraint unless π is greater than a critical value. We state this more formally as the following

[15] $\hat{c}$ is also increasing in δ.
[16] It is easily shown, additionally, that $\tilde{c}$ falls with increases in k and β.
[17] $\hat{c}$ is also increasing in δ, and decreasing in k and β.

Proposition 1. *Truth-telling is not an incentive-compatible strategy for all firms if* $\pi < \bar{\pi}$, *where* $\bar{\pi}$ *is defined by*

$$\bar{\pi} = \frac{\delta k - \delta(1-\beta)(\delta k - \bar{c})}{\bar{F} + \delta k - \delta(1-\beta)(\delta k - \bar{c})}$$

and has the property that $0 < \bar{\pi} < 1$.

Proof. $\bar{\pi}$ is the smallest value of π that makes the truth-telling constraint in Eq. (9) true for all values of $c \in [0, \bar{c}]$. Since the constraint is more likely to be satisfied when F_1 is smaller and F_2 is larger, $\bar{\pi}$ occurs at $F_1 = 0$ and $F_2 = \bar{F}$. Thus, for $\pi < \bar{\pi}$ it is impossible to adjust the fines to make truth-telling incentive-compatible for all firms. QED

Proposition 1 differs from the results of Kaplow and Shavell (1994) where truth-telling is an optimal strategy for all firms provided that $F_1 \leq \pi F_2$. The reason that our results differ is that we allow the regulator the power to exploit the information revealed by a firm that tells the truth. Specifically, the regulator's enforcement power allows it to force confessed violators to return to compliance. Since returning to compliance is costliest for the highest-cost firms, they are better off lying about their noncompliance status. As a result, truth-telling is not incentive compatible for all firms in our model.

It is helpful to see that if we set $\beta = 1$, our model reduces to the model of enforcement studied by Kaplow and Shavell (1994) and Malik (1993). Setting $\beta = 1$ reduces the truth-telling constraint in Eq. (9) to

$$F_1 + \delta k < \pi(F_2 + \delta k)$$

Interpreting $F_i + \delta k$ as the full penalty violators must pay makes this truth-telling constraint equivalent to that in Kaplow and Shavell (1994).

Why is the role of β so important? With $\beta = 1$, the firm begins each period with a broken PAD even following an order by the regulator to return to compliance. It is impossible, therefore, to force a firm to return to compliance and impossible, therefore, for information revealed in a confession of guilt to be used against the firm. As a result, truth-telling is incentive compatible for all firms. In contrast, when $0 < \beta < 1$, a guilty firm is forced back to compliance (with probability $1 - \beta$).

When $\pi < \bar{\pi}$, four general configurations for the payoffs can occur.

Case I. When both types of noncompliance strategies exist in equilibrium (truth telling and lie telling), there exists a $\tilde{c}$, shown in Fig. 1, such that all firms with $c < \tilde{c}$ will choose truth-telling, and those with $c > \tilde{c}$ will falsely report. Note that the proportion of firms that truthfully reveal their compliance status is

$$T = G(\tilde{c})$$

Since $\tilde{c}$ is decreasing in F_1, a reduction in F_1 necessarily increases truth-telling.

Case II. It is possible for only one noncompliance strategy (lie-telling) to exist in equilibrium. In this case case Z_{NT} lies everywhere above min$\{Z_{CT}, Z_{NF}\}$ as in Fig. 2. Here, firms following strategy S_{CT} are distinguished from those firms following strategy S_{NF} by c'.

The proportion of firms that truthfully report their compliance status in this case is

$$T = G(c')$$

Since c' is independent of F_1, truth-telling is unaffected by changes in F_1 in this case.

Case III. In this case Z_{CT} lies everywhere above min$\{Z_{NF}, Z_{NT}\}$ and so once again $\tilde{c}$ distinguishes the truth-tellers from those reporting falsely, but there are no fully complying firms. A reduction in F_1 increases truth-telling.

Case IV. If min$\{Z_{NT}, Z_{CT}\}$ is everywhere above Z_{NF} then there is no truth-telling in equilibrium.

When $\pi \geq \bar{\pi}$, there exist values for F_1 and F_2 that ensure that truth-telling gives the highest payoff for all firms. In this case, $T = 1$.

3.4. Compliance

Differences in PAD operating costs also allow us to identify those firms that comply. Of those firms that choose truth-telling, some will comply with the standards and some will not. The condition for compliance among truth-telling firms is that $Z_{CT} \leq Z_{NT}$. The equality defines a critical cost, $\hat{c}$, defined in Eq. (6) and shown in Fig. 1, such that firms with $c < \hat{c}$ comply and those with $c > \hat{c}$ violate. Such a $\hat{c}$ exists if $Z_{CT}|_{c=0} < Z_{NT}|_{c=0}$ or, alternatively, $F_1 > k(1-\delta)$.

The overall industry compliance rate depends on how many firms choose each of the three strategies. To illustrate, continue to assume that all three strategies exist (Case I). Then, the proportion of firms choosing always to comply is $G(\hat{c})$ – these firms have a 100% compliance rate. Next, the proportion of firms choosing to not comply but to report truthfully is $G(\tilde{c}) - G(\hat{c})$ – these firms have a compliance rate of $1-\beta$, the proportion of the time their PAD is operating. Finally, the proportion of firms choosing not to comply and to report falsely is $1 - G(\tilde{c})$ – these firms comply only when their machines are working but do not report the truth when their machines are not working. As a result, their noncompliance status is discovered only by inspections with probability π each period, after which they return to compliance. The steady-state proportion of these firms that have working

PAD machines, and therefore the compliance rate of these firms, can be shown to be $\pi(1-\beta)/[\beta+\pi(1-\beta)]$.[18]

The overall compliance rate is therefore

$$C = G(\hat{c}) + (1-\beta)[G(\tilde{c}) - G(\hat{c})] + \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)}[1-G(\tilde{c})]$$

Simplifying gives[19]

$$C = \beta G(\hat{c}) + \frac{\beta(1-\beta)(1-\pi)}{\beta+\pi(1-\beta)} G(\tilde{c}) + \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)} \tag{10}$$

It is easily shown that an increase in π and an increase in F_2 both increase the compliance rate. As in the standard Becker model, more frequent inspections and/or higher penalties for false compliance reports increases the compliance rate. However, the impact of F_1 on the compliance rate is unusual in this model and, indeed, stands as the contribution of this model towards a better understanding of the puzzling empirical phenomenon of high compliance rates and low expected fines. The role played by F_1 can be seen graphically by considering Fig. 1. As F_1 increases, the Z_{NT} curve shifts up, while the other curves remain unchanged. Thus, $\hat{c}$ increases, increasing the number of firms that comply always, and $\tilde{c}$ falls, reducing truth-telling. So, while the number of firms that comply voluntarily increases as F_1 increases, the number of violating firms that can be detected and forced back into compliance decreases. The first effect tends to increase the overall compliance rate, while the second effect tends to reduce it. In addition to these two

[18] To see this, denote by K_t the number of such firms that are in compliance at date t, and by M_t the number of firms out of compliance, and let $1=K_t+M_t$. The evolution of K_t is described by

$$K_t = (1-\beta)K_{t-1} + \pi(1-\beta)M_{t-1}$$

Since these firms never report when they are out of compliance and do not voluntarily fix broken PADs, firms that were in compliance in $t-1$ remain so only if the machine does not break down – with probability $1-\beta$. Those firms that were not in compliance in $t-1$ are inspected and forced into compliance with probability π and there is a probability $1-\beta$ that the PAD does not break down again.Using $M_{t-1}=1-K_{t-1}$ and setting $K_t=K_{t-1}$ for the steady state gives

$$K = \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)}$$

[19] By a similar logic, it follows that the compliance rate in Case II is

$$C = \beta G(c') + \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)}$$

which is increasing in F_2 and π. In Case III, it is

$$C = \frac{\beta(1-\beta)(1-\pi)}{\beta+\pi(1-\beta)} G(\tilde{c}) + \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)}$$

which is increasing in F_2 and π. Finally, in Case IV we have

$$C = \frac{\pi(1-\beta)}{\beta+\pi(1-\beta)}$$

which is independent of the fine structure, but increasing in π.

offsetting forces, the overall effect of a change in F_1 depends on the density of firms in the neighbourhoods of $\hat{c}$ and $\tilde{c}$. Before we consider this complication, we present our results initially using the assumption that the density of firms' costs is independent of the cost level. We relax this later.

Assumption 1. The cost structure of firms has a uniform density over $[0, \bar{c}]$. Denote this density simply by g, so that $G'(\cdot)=g>0$ for all $c\in(0, \bar{c})$.

We are able to establish the relationship between the overall compliance rate and the fine for violations. We state our result as follows:

Proposition 2. *Given Assumption 1 and $\pi<\bar{\pi}$, the compliance rate is a non-increasing function of the fine, F_1 and is a strictly decreasing function of the fine, F_1, if also $\tilde{c}>\hat{c}$.*

Proof. The impact of F_1 on the compliance rate is found by differentiating Eq. (10). Given Assumption 1, the proportion of firms choosing the lying strategy is greater than zero, and there are four cases to be considered.

Case I. When each of the strategies is optimal over some range of compliance costs, as shown in Fig. 1, we get

$$\frac{\partial C}{\partial F_1} = -\frac{\beta g(1-\delta)}{\delta[\pi(1-\beta)+\beta]} < 0$$

A reduction in F_1 unambiguously raises the compliance rate.

Case II. If strategy S_{NT} (violate but report truthfully) is dominated by either or both of the other two strategies, so that, as in Fig. 2, Z_{NT} lies everywhere above the envelope $\min\{Z_{CT}, Z_{NF}\}$, then

$$\frac{\partial C}{\partial F_1} = 0$$

F_1 has no impact on the compliance rate because no firms choose the strategy which involves paying the fine.

Case III. If strategy S_{CT} (full compliance) is dominated everywhere but both S_{NT} (noncompliance with truthful reporting) and S_{NF} (noncompliance and false reporting) are optimal, (which occurs when $F_1<k(1-\delta)$) then

$$\frac{\partial C}{\partial F_1} = -\frac{\beta[1-\delta(1-\beta)(1-\pi)]}{\delta[\pi(1-\beta)+\beta]}\, g < 0$$

In this case, a reduction in F_1 unambiguously raises the compliance rate.

Case IV. If Z_{NF} is everywhere below $\min\{Z_{CT}, Z_{NT}\}$, then

$$\frac{\partial C}{\partial F_1} = 0 \qquad \text{QED}$$

Proposition 2 challenges the conventional result in the theory of enforcement that compliance can be increased by increasing the fine for non-compliance. Instead, we obtain the result that the compliance rate may be increased by decreasing the fine for non-compliance. Why is this so? In Case I, in which all three strategies are present, lowering F_1 shifts Z_{NT} down in Fig. 1 but does not affect Z_{CT} or Z_{NF}. This lowers the proportion of firms that choose to comply ($\hat{c}$ in Fig. 1 falls). However, at the same time, the proportion of truth-telling firms increases. As a result, the proportion of firms that choose to violate but not report it decreases, ($\tilde{c}$ increases in Fig. 1). This has a positive effect on the industry compliance rate because those firms that now report their violations instead of filing false compliance reports can be returned to compliance sooner, since they are identified immediately instead of with probability π per period. Because of the relative slopes of Z_{CT} and Z_{NF} and with a uniform density, more violating firms become truth-tellers than do complying firms become violators. For a given reduction in F_1, $\tilde{c}$ is increased by more than $\hat{c}$ is reduced.

In Case III, in which the strategy of full compliance does not exist, lowering F_1 cannot reduce the proportion of firms that comply because it is already zero. As a result we only have the second effect in which untruthful firms switch to being truthful as $\tilde{c}$ increases.

Finally, we note as a corollary to Proposition 1 that, as π approaches $\bar{\pi}$, then the compliance rate, C, approaches $1-\beta$. This can be seen readily from Eq. (10) since as π approaches $\bar{\pi}$, (and with $F_1=0$, $F_2=\bar{F}$), $\hat{c}=-k(1-\delta)$ so $G(\hat{c})=0$, and $\tilde{c}$ approaches $\bar{c}$ so $G(\tilde{c})=1$ (all firms choose the strategy of noncompliance with truth-telling). If $\pi>\bar{\pi}$, a higher compliance rate can be achieved. However, if the truth-telling constraint in Eq. (9) holds with equality, an increase in F_1 will again lower the compliance rate because it will lead to less than complete truth-telling in which case the results of the proof of Proposition 2, Case I or Case III apply.

4. Cost-minimizing enforcement

We assume that the agency's objective is to minimize its cost of enforcement subject to achieving a given compliance rate.[20] The agency achieves this by

[20] Alternatively, we could assume the agency has a given budget and chooses the parameters to maximize the compliance rate. Our results are the same in either case.

choosing an enforcement policy $\{F_1, F_2, \pi\}$, with $0 \leq F_1 < F_2 \leq \bar{F}$ and $\pi \in (0, 1)$.[21] The exogenous minimum compliance rate is C and is assumed to be feasible for at least some $\{F_1, F_2, \pi\}$.

Let N be the number of firms reporting compliance. Then the number of firms that must be inspected each period is πN and the cost is $w\pi N$, where w is the cost of inspection. Note that the exact equations for C and N depends on how many strategies are in effect. Equations for C have been given earlier, while equations for N are readily calculated. In Case I, for example, N is the sum of the measure of firms reporting compliance truthfully ($G(\hat{c})$) and the measure of firms reporting falsely that they are in compliance $(1-G(\tilde{c}))$,[22] or

$$N = G(\hat{c}) + 1 - G(\tilde{c})$$

In addition to inspection costs, we assume the agency incurs an enforcement cost, f, for each firm that is forced back into compliance.[23] All of the NT firms that experience PAD failure (a fraction β) are forced back into compliance resulting in an enforcement cost of $f\beta(G(\tilde{c})-G(\hat{c}))$. In addition, a fraction π (the inspection rate) of the noncompliant NF firms is forced back into compliance, resulting in an enforcement cost of $f\pi\beta/(\beta+\pi(1-\beta)(1-G(\tilde{c}))$.The enforcement agency's problem is

$$\min_{F_1,F_2,\pi} w\pi[G(\hat{c}) + 1 - G(\tilde{c})] + f\beta[G(\tilde{c}) - G(\hat{c})] + \frac{f\beta\pi}{\beta + \pi(1-\beta)}[1 - G(\tilde{c})] \tag{P}$$

subject to

$$C(F_1, F_2, \pi) \geq C^*$$

and $0 \leq F_1 < F_2 \leq \{\bar{F}\}$, $\pi \in (0, 1)$.

The problem in Eq. (P) may be solved in two stages. Firstly, the regulator chooses a fine structure for a *given* π. Then, given the optimal fine structure, the regulator chooses π. The outcome of the first stage is an optimal fine structure as a function of π. The key to this relationship is the size of the compliance rate target relative to $1-\beta$. The following result applies to the case where $C^* < 1-\beta$

[21] We are implicitly assuming that the setting of the emissions constraint is exogenous to the enforcement agency. The results in Harford and Harrington (1991) imply that if the agency can control both the compliance rate and the enforcement policy, the social cost of control can be lowered by relaxing the compliance rate somewhat.

[22] In Case II and Case IV we have $N=1$, while in Case II $N=1-G(\tilde{c})$.

[23] We are grateful to an anonymous referee for this suggestion.

Proposition 3. *Given Assumption 1, and* $C^* < 1 - \beta$, *the solution to the cost minimization problem* Eq. (P) *requires setting* $F_1 = 0$, $F_2 = \bar{F}$, *and* $0 < \pi < \bar{\pi}$ *if* $w\pi > \beta f$.

The proof is straightforward because the expected cost is strictly increasing in F_1 if $w\pi > \beta f$ in Cases I and III. This, combined with the result already obtained that the compliance rate is non-decreasing in F_1 means cost is minimized by setting F_1 to its minimum value. In Case IV, both expected cost and the compliance rate are unaffected by the value of F_1; setting $F_1 = 0$ is arbitrary in this case. In all four cases, expected cost is decreasing and compliance increasing in F_2 which implies setting $F_2 = \bar{F}$ to minimize cost.

For completeness, we are able to state the conditions for the optimal choice of π. In view of Proposition 3, the problem is:

$$\min_{\pi} w\pi N(0, \bar{F}, \pi)$$

subject to $C^* \leq C(0, \bar{F}, \pi)$.

Proposition 3 indicates that a given level of compliance is achieved at least cost by setting F_1, the fine for noncompliance, equal to zero if $w\pi > \beta f$ and F_2, the fine for false reporting, as high as possible. These results reflect the fact that the solution involves making the difference between F_2 and F_1 as large as possible in order to maximize the proportion of firms that truthfully report violations. This both decreases the number of firms that need to be inspected and increases the number of violators that are identified (through self-reporting) and returned to compliance rapidly.

Why is Proposition 3 conditional on $w\pi > \beta f$? Intuitively, lowering F_1 shifts some firms from the voluntary compliance strategy (CT) to the non-compliance but truth-telling strategy (NT). These firms no longer need to be inspected and this results in inspection cost savings equal to $w\pi$. However, these firms now have to be forced to return to compliance when their PADs fail, so there are additional enforcement costs equal to βf. If the inspection cost savings outweigh the additional enforcement costs, it is optimal to set $F_1 = 0$ to take full advantage of the net savings.[24] Indeed, if it were possible to set $F_1 < 0$ in Eq. (P) it would be optimal to do so, implying that truth-telling is desirable in this model.[25]

It is quite plausible that inspection costs are large relative to enforcement costs for some types of pollutants and industrial processes. For example, inspections normally require a site visit whereas an enforcement action may be as administra-

[24] There are actually additional cost savings because lowering F_1 shifts some noncompliant firms from the false reporting to the truthful reporting strategy which further reduces the need for inspections.

[25] We are forcing a non-negativity condition on F_1. This reflects asymmetries that usually exist in the implementation of incentive mechanisms and taxation. Subsidies for polluting firms (reporting truthfully) may be politically infeasible and may also have obvious moral hazard properties.

tively straightforward as a legally enforceable written order to return to compliance. At the same time, however, the model predicts that in situations where the enforcement cost, f is high and/or the frequency of equipment failure, β is high, (so that enforcement actions would be required) a cost-minimizing policy is more likely to involve $F_1 > 0$ to encourage voluntary return to compliance.

We do not model the regulator's objective as the minimization of pollution damage costs plus enforcement and compliance costs. Thus, unlike Kaplow and Shavell (1994), we have not attempted to model the socially optimal enforcement policy. Instead, our interest is in showing that the observation of low expected penalties for noncompliance is consistent with cost-minimizing behaviour on the part of regulators charged with the task of achieving a given compliance rate.[26] However, it is interesting to note that in the case of an extremely high exogenous compliance target (as it would be presumably for pollutants such as radioactive wastes that create very high social damages), the cost-minimization problem Eq. (P) leads to a qualitatively different solution. We state this formally as

Proposition 4. *If $C^* > 1 - \beta$, then, given Assumption 1 and $w\underline{\pi} > f\beta$, the solution to the cost-minimization problem* Eq. (P) *requires setting $F_2 = \bar{F}$, $\pi > \bar{\pi}$ and $F_1 > 0$ to satisfy*

$$\delta k < F_1 + \delta k < \pi(\bar{F} + \delta k)$$

Proof. Suppose $C^* = 1 - \beta$. If $F_1 = 0$ and $F_2 = \bar{F}$, then it is necessary to set $\pi = \bar{\pi}$ to achieve this compliance rate. Since $\bar{\pi}$ is the smallest π which induces truth-telling for all firms, then $F_1 = 0$, $F_2 = \bar{F}$ and $\pi = \bar{\pi}$ minimizes cost. Now suppose $C^* > 1 - \beta$. Raising F_1 alone will only decrease the compliance rate (by Proposition 2). Likewise, raising π alone will not raise the compliance rate because no firms are filing false reports. Raise F_1 and π together so as to keep the truth-telling constraint just binding raises the compliance rate and it does so at least cost by minimizing the amount by which π must be increased. Low cost firms are now encouraged to comply voluntarily by raising F_1 but high-cost firms are discouraged from filing false reports by raising π by just enough to make $\tilde{c} = \bar{c}$. Setting $\tilde{c} = \bar{c}$ gives the expression

$$F_1 + \delta k = \frac{\pi(\bar{F} + \delta k) - \alpha\bar{c}}{1 - \alpha}$$

which F_1 and π must satisfy. This,determines F_1 and π. It is straightforward that as $\pi \to 1$, $F_1 + \delta k \to \bar{F} + \delta k$ and that for $\pi < 1$, $F_1 + \delta k < \pi\ (\bar{F} + \delta K)$. QED

[26] As a result of this difference in objective functions, F_1 and F_2 in our model do not have the same Pigouvian-tax kind of interpretation as they do in Kaplow and Shavell's model.

5. Extensions

The main part of this paper has worked through a model of environmental regulation which has been restricted in two principal ways. First, we ruled out the possibility that firms might voluntarily disable the PAD so as to save on running costs. This is relaxed in the following subsection. Secondly, we assumed that running costs were uniformly distributed across firms. We re-examine this assumption and discuss alternative sufficient conditions for the main results of the paper to hold. However, in this section we consider only the case in which truth-telling is not incentive-compatible for all firms.

5.1. Initiating noncompliance

Our concern so far has been with the consequences of noncompliance being forced on firms through the (random) breakdown of the abatement technology. In this section we consider briefly an additional option for a firm – that of deliberately disconnecting the PAD. In this case the firm is voluntarily initiating the state of noncompliance and is not simply a victim of the exogenous breakdown of the equipment.

The possibility of deliberate noncompliance substantially weakens the enforcement power we have given to the regulator in our model. Although a firm that adopts this noncompliance strategy can be ordered to return to compliance, such an order is ineffective because the firm immediately returns to a state of noncompliance. The regulator therefore has no power to prevent a firm from always being out of compliance. This severely constrains the extent to which compliance gains can be achieved through lower fines for noncompliance as was the case in the previous section.[27]

This extension introduces two new potential courses of action for the firm in addition to those discussed earlier. The firm may choose to disconnect the PAD and then to report that it has done so, or it may report falsely. Denote by S_{DT} the strategy of disconnecting the PAD and truthfully reporting it and Z_{DT} the associated expected cost of this strategy. Similarly, let S_{DF} denote the strategy of disconnecting the PAD and filing a false report on compliance status, and let Z_{DF} be the resulting expected cost to the firm. We assume that technological failure does not befall a deliberately disconnected PAD and that any disconnection brought to the attention of the regulator, either by a truthful report or by random inspection, requires the firm to restart the device at a cost $s > 0$. The fine structure is the same as before. It is straightforward to see that Z_{DT} is given recursively by

[27] A more sophisticated enforcement power (such as an increasing fine structure for repeat offenders) is clearly required to deal with deliberate noncompliance. It is beyond the scope of this paper to consider different forms of enforcement power.

$$Z_{DT} = F_1 + \delta s + \delta Z_{DT}$$

where the discount term on s indicates that the firm is given one period to complete the start-up. Without loss of generality, we assume that the firm can immediately disconnect the device next period.

Similarly, a firm may choose to disconnect and file a false report. It is easy to see that Z_{DF} in this case is given by

$$Z_{DF} = \pi(F_2 + \delta s + \delta Z_{DF}) + (1 - \pi)\delta Z_{DF}$$

If the firm is inspected, with probability π, it is fined and forced back into compliance, thus incurring the start-up cost. However, the firm immediately disables the PAD and returns to noncompliance.

Notice that the running cost is avoided under both strategies, so that only the lower-cost of the two strategies would ever be chosen by any firm. We consider first the case in which $Z_{DF} < Z_{DT}$; that is, false-reporting dominates truthful reporting of the disconnection of the PAD. This case occurs if the expected penalty from false reporting, $\pi(F_2 + \delta s)$, is less than the certain penalty from truthful reporting, $F_1 + \delta s$. This condition is satisfied even with $F_1 = 0$ and $F_2 = \bar{F}$ as long as

$$\frac{\pi}{1 - \pi}\bar{F} < \delta s$$

which requires that the start-up cost, s, be sufficiently large. Intuitively, it is unprofitable to truthfully report disconnecting the PAD in this case because of the high cost of re-starting the PAD each period.

In this case, the results concerning the aggregate compliance rate and the cost-minimizing enforcement policy are qualitatively unchanged:

Proposition 5. *If, in addition to the noncompliance strategies already examined, noncompliance can be initiated by deliberate disconnection of the PAD, but false reporting dominates truthful reporting of this strategy then, (a) the aggregate compliance rate is a non-increasing function of F_1 and (b) the cost-minimization problem* Eq. (P) *is solved by setting $F_1 = 0$, $F_2 = \bar{F}$, and* $\pi \in (0, 1)$ *given Assumptions 1 and 2.*

The proof is presented in Appendix A.

Next, we consider the case in which $Z_{DF} > Z_{DT}$; that is, truthful reporting dominates false reporting of the disconnection of the PAD. This case is more likely to occur when start-up costs, s, are small and necessarily occurs if $s = 0$. A firm following this strategy deliberately and openly violates the pollution standard repeatedly, and willingly pays the fine F_1 and the discounted start-up costs, δs. In this case, we get

Proposition 6. *If, in addition to the noncompliance strategies already examined, noncompliance can be initiated by deliberate disconnection of the PAD, and truthful reporting dominates false reporting of this strategy then, (a) the aggregate compliance rate is an ambiguous function of* F_1 *and (b) the cost-minimization problem* Eq. (P) *is, in general, solved with* $F_1>0$ *and* $F_2=\overline{F}$, $\pi\in(0, 1)$.

Again, the proof is in Appendix A.

Propositions 5 and 6 show that when enforcement power is weakened by the ability of polluters to openly defy enforcement orders, the ability of the regulator to obtain high compliance rates with low fines hinges on the technical ability of firms to switch back and forth between states of compliance and noncompliance. If it is sufficiently costly to move between these states, firms that choose to deliberately disconnect the PAD will continue to (falsely) report compliance to the agency. In this case, Proposition 4 shows that the cost-minimizing enforcement policy involves setting the noncompliance fine equal to zero under Assumptions 1 and 2. On the other hand, if the cost of moving between the states of compliance and noncompliance is sufficiently low that some firms choose to deliberately disconnect the PAD and truthfully report that they have done so, it is impossible to obtain higher compliance rates by lowering the noncompliance fine.

5.2. Sufficient conditions

The focus of the paper has been on a leading special case in which the distribution of costs is assumed to be uniform. This has permitted us to ignore the influence of the cost structure on firm behaviour in response to changes in the fine structure. To consider the general case, we can look at the derivative of Eq. (10) with respect to F_1 for an arbitrary, but differentiable distribution function. A simple calculation shows this to be

$$\frac{\partial C}{\partial F_1}=\beta\left([g(\hat{c})-g(\tilde{c})]-\frac{(1-\beta)(1-\pi)(1-\delta)}{\alpha[\beta+\pi(1-\beta)]}\,g(\tilde{c})\right)$$

A large number of alternative sufficient conditions for C to be decreasing in F_1 are now apparent; for example, that g always has more mass around $\tilde{c}$ than around $\hat{c}$. In general, the intuitive interpretation of any alternative to the uniform density as a sufficient condition, is that C will fall with F_1 as long as not ‘‘too many’’ additional firms switch from being compliers to being noncompliers relative to the number switching from being truth-tellers to being false-reporters for each dollar reduction in F_1.

6. Conclusion

The point of departure for this paper is the observation that the predictions of the standard theory of enforcement are at odds with a stylized fact of pollution-standards enforcement: compliance rates are high even though expected penalties for violations are low. The purpose of this paper is to argue that this observation is in fact consistent with theoretical predictions when the theory is extended to include a self-reporting requirement and enforcement. Both of these elements are present in practice. Our main findings are that in this extended enforcement model it is possible, under some conditions, to achieve higher compliance rates with lower fines (the reverse of the conventional prediction) and that, when this is true, the cost of enforcing a given level of aggregate pollution is minimized by setting the fine for noncompliance equal to zero.

Lowering the fine for noncompliance has two effects in our model. First is the conventional effect that lowering the fine reduces the proportion of firms that choose to comply always. This reduces the compliance rate. Second is the new effect that lowering the fine raises the proportion of noncomplying firms that file truthful reports about their compliance status. This increases the overall compliance rate because it leads to earlier detection of noncompliers. That is, rather than having to rely on identification through random inspections, the regulator, by lowering the fine for noncompliance, can rely on noncompliers to identify themselves. This permits the regulator to take actions to order these firms to return to compliance sooner than would otherwise be the case.

We found the conditions under which the second effect dominates the first to be plausible. A necessary condition is that it not be too easy for firms to alternate between states of compliance and noncompliance. If the cost of moving between states is sufficiently low and the fine is also low, the optimal strategy for most firms is to never comply but to report truthfully. The only way to prevent this in our model is to raise the fine for noncompliance. The model predicts then, that in situations where the cost of alternating between compliance and noncompliance is low (as may be the case for illegal waste disposal and, to suggest a non-pollution example, speeding infractions), enforcement policies would include substantial fines for noncompliance.

If the cost of alternating between states of compliance and noncompliance is sufficiently high, as it would be if the control method is an integral part of the production technology or the process of disabling and re-enabling the pollution control equipment is costly (as it might be for catalytic converters and flue gas scrubbers, for example), our results depend on the density of firms at low versus high cost of abatement. In the case of the uniform density, we find that the aggregate compliance rate is non-increasing in the fine for noncompliance and that setting the fine equal to zero minimizes the cost of enforcing a given level of aggregate pollution emissions. These results hold for other distributions as long as not too many firms adopt noncompliance strategies and a sufficiently large number adopt truth-telling behaviour.

Acknowledgements

We are grateful to Andrew Keeler, Charles Kolstad, Ross McKitrick, two referees, and seminar participants at the following universities: Guelph, Munich; Wilfrid Laurier, York, Waterloo, Victoria, Brock and Cornell. An earlier version of this paper was presented at the fourth annual conference of the Canadian Resource and Environmental Study Group, and the eighth annual conference of the European Association of Environmental Economists in Tilburg.

Appendix A

Proof of Proposition 5. To prove part (a), we show that the results in the proof of Proposition 2 are not changed by the addition of the new strategy. This, together with the demonstration that the number of firms to inspect, wN, is non-decreasing in F_1 proves part (b).

When the new strategy is present, the number of cases that can arise increases. However, many of these cases exclude the truth-telling strategy, S_{NT}. In these cases, the aggregate compliance rate and the number of firms to inspect are both independent of F_1 because no firms choose the strategy that results in this fine being imposed. As before, the value of F_1 is irrelevant in these cases so we arbitrarily set it equal to zero.

Four cases can arise that include the truth-telling strategy, S_{NT}, and therefore involve the fine F_1 being imposed. We consider each in turn.

First, all four strategies may co-exist in equilibrium: (S_{CT}, S_{NT}, S_{NF}, S_{DF}). The compliance rate is

$$C = G(\hat{c}) + (1-\beta)[G(\tilde{c}) - G(\hat{c})] + \gamma[G(c_3) - G(\tilde{c})]$$

where $\gamma = \pi(1-\beta)/(\beta + \pi(1-\beta))$ and c_3 is the solution to $Z_{DF} = Z_{NF}$. In this case, $\partial c_3/\partial F_1 = 0$. Hence, $\partial C/\partial F_1 < 0$ as under Assumption 1 and $\pi < \bar{\pi}$. Next, the agency's expected cost is

$$w\pi[G(\hat{c}) + 1 - G(\tilde{c})] + f\beta[G(\tilde{c}) - G(\hat{c})] + \frac{f\beta\pi}{\beta + \pi(1-\beta)}[G(c_3) - G(\tilde{c})] + f\pi[1 - G(c_3)]$$

which is increasing in F_1.

Second, the S_{NF} strategy is dominated everywhere but the other three remain: (S_{CT}, S_{NT}, S_{DF}). The compliance rate is

$$C = G(\hat{c}) + (1-\beta)[G(c_3) - G(\hat{c})]$$

where c_3 is the solution to $Z_{NT} = Z_{DF}$. In this case, $\partial c_3/\partial F_1 < 0$ but $\partial C/\partial F_1 = 0$ under Assumption 1 and $\pi < \bar{\pi}$. Next, the agency's expected cost is

$$w\pi[G(\hat{c}) + 1 - G(c_3)] + f\beta[G(c_3) - G(\hat{c})] + f\pi[1 - G(c_3)]$$

which is increasing in F_1.

Third, the S_{CT} strategy is dominated but the other three exist in equilibrium: (S_{NT}, S_{NF}, S_{DF}). The compliance rate is

$$C = (1 - \beta)G(\tilde{c}) + \gamma[G(c_3) - G(\tilde{c})]$$

where c_3 is the solution to $Z_{NF} = Z_{DF}$. In this case, $\partial c_3/\partial F_1 = 0$ and $\partial C/\partial F_1 < 0$ under Assumption 1 and $\pi < \bar{\pi}$. In addition, the agency's expected cost is

$$w\pi[1 - G(\tilde{c})] + f\beta G(\tilde{c}) + \frac{f\beta\pi}{\beta + \pi(1 - \beta)}[G(c_3) - G(\tilde{c})] + f\pi[1 - G(c_3)]$$

which is increasing in F_1.

Fourth, only two strategies exist in equilibrium: (S_{NT}, S_{DF}). The compliance rate is

$$C = (1 - \beta)G(c_3)$$

where c_3 is the solution to $Z_{NT} = Z_{DF}$. In this case, $\partial c_3/\partial F_1 < 0$ which makes $\partial C/\partial F_1 < 0$ under Assumption 1 and $\pi < \bar{\pi}$. In addition, the agency's expected cost is

$$w\pi[1 - G(c_3)] + f\beta G(c_3) + f\pi[1 - G(c_3)]$$

which is increasing in F_1.

Proof of Proposition 6. The proof is straightforward when $s = 0$ for in this case, setting $F_1 = 0$ would induce all firms to adopt the strategy of disconnection combined with truthful reporting; in that case, the aggregate compliance rate would be zero. For larger values of s, the proof follows the same steps as in Proposition 4 with four cases to examine. The expressions for C and expected cost are identical to those in the proof of Proposition 5. The difference is that c_3 is determined by the conditions $Z_{NF} = Z_{DT}$, $Z_{NT} = Z_{DT}$, $Z_{NF} = Z_{DT}$, and $Z_{NT} = Z_{DT}$ respectively. In each case, $\partial c_3/\partial F_1 > 0$ which is enough to make the aggregate compliance rate sometimes increasing and sometimes decreasing in F_1 under Assumption 1 and $\pi < \bar{\pi}$, depending on parameter values. In the other cases, C is independent of F_1. This proves part (a) of the proposition. Part (b) follows directly.

References

Alm, J., McClelland, G.H., Schutze, W.D., 1992. Why do people pay taxes?. Journal of Public Economics 48, 21–38.

Becker, G.S., 1968. Crime and punishment: an economic approach. Journal of Political Economy 76, 169–217.

Cropper, M.L., Oates, W., 1992. Environmental economics: a survey. Journal of Economic Literature 30 (2), 675–740.

Harford, J.D., 1978. Firm behavior under imperfectly enforceable pollution standards and taxes. Journal of Environmental Economics and Management 5, 26–43.

Harford, J.D., 1987. Self-reporting of pollution and the firm's behavior under imperfectly enforceable regulations. Journal of Environmental Economics and Management 14, 293–303.

Harford, J.D., Harrington, W., 1991. A reconsideration of enforcement leverage when penalties are restricted. Journal of Public Economics 45, 391–395.

Harrington, W., 1988. Enforcement leverage when penalties are restricted. Journal of Public Economics 37, 29–53.

Kaplow, L., Shavell, S., 1994. Optimal law enforcement with self-reporting of behavior. Journal of Political Economy 102, 583–606.

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

Ontario Ministry of the Environment, 1986–89. Annual direct discharge reports, Toronto, ON.

Russell, C.S., Harrington, W., Vaughan, W.J., 1986. Enforcing pollution control laws. Resources for the Future, Washington, D.C.

Storey, D.J., McCabe, P.J., 1980. The criminal waste disaster. The Scottish Journal of Political Economy 27, 30–40.

Swierzbinski, J.E., 1994. Guilty until proven innocent-regulation with costly and limited enforcement. Journal of Environmental Economics and Management 27, 163–184.