JOURNAL OF ENVIRONMENTAL ECONOMICS AND MANAGEMENT 27, 127–146 (1994)

Guilty until Proven Innocent — Regulation with Costly and Limited Enforcement[1]

JOSEPH E. SWIERZBINSKI

Department of Economics and School of Natural Resources and the Environment, University of Michigan, Ann Arbor, Michigan 48109

Received April 16, 1993; revised July 2, 1993

We consider optimal regulations for a polluting firm when regulators cannot observe emission control costs and can only observe emissions via costly monitoring. Fines (or subsidies) for enforcing compliance are also limited. The optimal regulations resemble a deposit–refund system. The firm reports its emissions and pays an initial tax based on this report. If the firm is monitored, it receives a rebate when actual and reported emissions coincide. The enforcement constraints and the firm's rights determine whether the incentives to reduce emissions are optimally provided by varying the rebate for compliance, the monitoring probability, or the initial tax.

1. INTRODUCTION

Emission or effluent fees are an important part of the system for regulating pollution in several European countries. In textbook discussions, the most important role for such fees, or Pigouvian taxes, is to provide polluters with the correct incentives by internalizing external costs.[2] However, a number of economists have questioned the extent to which several European fee systems actually influence the incentives of polluters.[3] These economists observe that emission fees have typically been set at low levels and that a polluter's payment often depends on an estimated emission level rather than on the polluter's actual emission of pollutants. If a polluter's payment never depends on the amount he pollutes, then the payment obviously provides no incentive to reduce the level of pollution.

European fee systems are actually more complicated. In some systems, some polluters pay a fee based on their actual emissions while others pay a fee that is unrelated to the amount they pollute. It is also noteworthy that, in at least one case, a polluter may request to be monitored and have his estimated fee adjusted if his actual emission of pollutants differs from his estimated emission.[4] This raises the interesting possibility that an incentive to reduce pollution may be provided not by an "up-front" fee, as in the textbook analysis, but rather by variations in a "rebate" which polluters may receive when their actual emissions are monitored.

[1]I am grateful to Ken Binmore, Charles Kolstad, Steve Salant, an anonymous referee, and participants in the 1991 Canadian Resource and Environmental Economics Study Group, the 1992 Summer Institute in Game Theory at Stonybrook, and a 1992 conference on economic theory at the University of Michigan for helpful comments and suggestions. Of course, any errors are my own.

[2]See, for example, Baumol and Oates [6].

[3]Several recent surveys of environmental regulation take this view. See p. 692 of Cropper and Oates [18] and pp. 104–105 of Hahn [24].

[4]See Bower *et al.* [13, pp. 128–129] and OECD [35, p. 25].

0095-0696/94 $6.00

How should the incentives to reduce pollution best be provided? We consider the optimal regulation of polluting firms in the case where (1) it is costly to monitor the actual emissions of polluters, (2) the regulator is limited in the extent to which he can reward compliance or punish noncompliance, and (3) the regulator, although he knows the distribution of abatement costs, cannot observe the costs of any individual firm. These assumptions appear to be consistent with the characteristics of at least some systems for enforcing pollution controls laws.[5]

Given these limitations on the regulator's information and enforcement power, the following framework provides a reasonable description of the possibilities for regulating pollution via emission fees.[6] (1) A polluter reports what his level of emission will be and pays an initial fee or tax that depends only on his report. (2) The regulator then monitors the pollution level of the firm with a prespecified probability that typically depends on the polluter's reported emission level. For simplicity, we assume that the regulator's monitoring technology is perfectly accurate although costly. (3) If the polluter's emissions are monitored, then the regulator is also precommitted to punish the polluter with a fine or to reward him with a rebate depending on whether the polluter's actual emissions are consistent with his initial report.

We consider two cases. In the bulk of the paper, we assume that the regulator can commit to monitoring a firm with any probability. We show that in this case a random monitoring policy is generally optimal. That is, a firm is typically monitored with a probability that is strictly between 0 and 1. A firm that reports a lower level of emission is more likely to be monitored. As well as levying fines for noncompliance, it is optimal to offer a rebate to a firm that is monitored and found to be producing its reported level of emission. If the maximum possible fine for noncompliance is small compared to the maximum rebate, then almost all of the incentive to reduce emissions is provided by variations in the probability that a firm is monitored. At the other extreme, if the maximum rebate that the regulator is permitted to offer is small compared to the maximum fine, then almost all of the incentive to reduce emissions comes from variations in the initial tax.

We also consider the case where the regulator's ability to commit to a monitoring policy is constrained by the firm's right to be monitored. This is an important issue, since a firm that expects a rebate for compliance will wish to be monitored. One might expect that rebates would not be useful in such circumstances since they encourage firms to demand that the regulator engage in costly monitoring. However, this intuition is only correct if the maximum fine is large enough to provide effective deterrence. If the maximum fine is small, then it will be optimal to provide rebates for at least some reported levels of emission. In the extreme but not implausible case where the maximum fine is very small, all types of firm receive rebates, a firm is always monitored, and, in one scenario, all firms pay the same initial tax and all of the incentive to reduce emissions is provided by variations in the size of the rebate.

[5]Russell [38] observes that the regimes for enforcing pollution control laws in the United States and the United Kingdom are characterized by: (i) a heavy reliance on self-reporting, (ii) infrequent monitoring of polluters, and (iii) low penalties for violations.

[6]Given the structural assumptions of our model, a regulator can achieve at least as large a payoff by using the regulatory framework in our paper as he can by using any of a large class of alternative regulations. This point is discussed further in Section 2 of the paper.

The formal model developed in the paper adopts the approach of Baron and Myerson [5], who consider a regulator that wishes to control the output of a single firm without being fully informed about the firm's costs.[7] In the context of pollution control, a firm's "output" can be interpreted as its reduction in emissions from some base level. However, unlike the present paper, Baron and Myerson's model (and most of the literature on regulation under uncertainty) assumes that the regulator can costlessly observe the firm's output and impose an arbitrary, nonlinear tax as a function of the firm's actual output subject only to participation and incentive compatibility constraints.[8] Note that since there is only one regulatory instrument, the nonlinear tax, in the Baron–Myerson model, the question of which combination of instruments is best suited to provide the incentives for pollution reduction cannot be addressed.

In terms of methodology, the literature on income taxation and auditing, including papers by Border and Sobel [12], Chander and Wilde [16], Melumad and Mookherjee [32], Mookherjee and Png [33], and Reinganum and Wilde [37] is closest to the current paper. In most of this literature, the taxpayer has no choice to make, other than reporting an income level. Hence, the problem addressed in our paper of motivating an agent to both signal information about costs (via the reported emission level) and choose an action appropriately does not arise.[9]

Another literature related to our paper extends Baron and Myerson's [5] model to allow for audits of the firm's production costs. Papers in this literature include Baron [2], Baron and Besanko [4], Chapter 4 of Besanko and Sappington [9], and Demski *et al.* [20]. These papers follow Baron and Myerson in assuming that the regulator can costlessly observe the firm's output. Hence, the opportunity to audit the firm's production costs provides the regulator in these models with the chance to obtain information in addition to that available to the Baron–Myerson regulator. In contrast, the regulator in our model never knows more than the Baron–Myerson regulator and only obtains information that the Baron–Myerson regulator gets for free by engaging in costly monitoring. As a result, the constraints

[7]Baron [3], Besanko and Sappington [9], and Caillaud *et al.* [14] survey the large literature on regulated firms that has developed using the Baron–Myerson approach. Dasgupta *et al.* [19] and Spulber [42] develop models of environmental regulation that are similar to Baron and Myerson's model.

[8]Caillaud *et al.* [15] consider a model where a principal can impose a tax based only on a noisy observation of an agent's output. However, the principal in this model cannot pay for better information about the agent's output, and, consequently, many of the issues considered here, such as the optimal monitoring strategy and the optimal tax/rebate once an agent is monitored, do not arise. Ortuno-Ortin [36] develops a model of pollution control where a firm has private information about its pollution control costs and the regulator must engage in a costly inspection to discover the firm's level of emissions. The focus of Ortuno-Ortin's paper is quite different from that of the current paper; for example, rebates for compliance are not allowed in Ortuno-Ortin's model.

[9]Mookherjee and Png [33] consider a model where initially identical taxpayers have income levels that depend on an unobservable choice of effort. Unlike the firm in our model, Mookherjee and Png's agents do not have private information when they choose their effort levels. Moreover, by making structural assumptions which are common in the regulation literature but not in the literature on income taxation, we are able to obtain sharper results concerning the optimal regulations than those obtained by Mookherjee and Png. Other papers that combine auditing and principal–agent models include Baiman and Demski [1], Dye [22], and Evans [23]. The first two papers develop models where the principal costlessly observes output and can, in addition, pay to measure the agent's effort. Evans [23] develops a model where the principal observes nothing and receives a fixed payment unless the agent calls for an audit of both output and effort. As with Mookherjee and Png, the agents in these models do not have private information when they choose their effort levels.

on the optimal regulations are quite different in the two types of models, and even the qualitative properties of the optimal regulations differ. For example, in the limiting case where auditing or monitoring is too costly to be feasible, the Baron–Besanko regulator can still control the firm's output using Baron and Myerson's optimal tax policy. On the other hand, the only options available to the regulator in our model when monitoring is not possible are to shut the firm down with a prohibitive tax or to allow the firm to operate, in which case all types of firm will produce the lowest output or, equivalently, emit the highest level of pollutants.[10] As another example of the differences between the two types of model, in the Baron–Besanko framework it is optimal to audit firms with high costs more frequently, while in our model it is optimal to monitor firms with low costs more frequently.

When an up-front payment is combined with a rebate for compliance, the resulting regulation resembles a deposit–refund system or performance bond. A number of authors have considered the use of deposit–refund systems to control environmental problems.[11] Costanza and Perrings [17] and Cropper and Oates [18] observe that deposit–refund systems may be useful under conditions of uncertainty because they shift the "burden of proof" from the regulator to the polluter. What has been lacking until now is a demonstration that a regulatory system with deposit–refund features is optimal in the context of regulation under uncertainty and an analysis of how such a system interacts with other aspects of regulatory policy, such as the monitoring decision.[12]

The paper proceeds as follows. Section 2 describes the structural assumptions of the model and sets up the optimization that determines the optimal regulations for the case where the regulator can commit to monitoring with any probability. Section 3 discusses these optimal regulations. Section 4 discusses the optimal regulations for the case where the regulator's ability to commit to a monitoring policy is limited by the firm's right to be monitored, and Section 5 contains some concluding remarks.

2. THE MODEL

This section introduces assumptions and notation and describes the model analyzed in the rest of the paper.

Consider a risk-neutral regulator who wishes to control the output of a single firm which has private information about its production costs. In addition to uncertainty about the firm's costs, the regulator must monitor the firm in order to

[10] In the opposite limiting case where the cost of monitoring or auditing approaches zero, the audit is perfectly accurate, and the maximum penalties for misreporting approach infinity, the Baron–Besanko regulator achieves the expected net benefits of a perfectly informed regulator who observes both cost and output, and the regulator in our model achieves only the expected net benefits of a Baron–Myerson regulator.

[11] See Bohm [10], Bohm and Russell [11], Costanza and Perrings [17], and Solow [40] for discussions of deposit–refund systems in environmental contexts.

[12] A number of papers including Beavis and Dobbs [7], Beavis and Walker [8], Downing and Watson [21], Harford [25, 26], Harrington [27], Linder and McBride [28], Malik [29, 30], Martin [31], and Seegerson [39] consider issues related to the enforcement of environmental regulations. An important difference between these papers and the current paper is that the current paper allows the firm to have private information about its emission control costs while those of the above papers which consider the design of optimal regulations assume that the regulator and the firm are equally well informed about the costs of controlling the firm's emissions.

observe and, if necessary, penalize deviations from the desired output level. Monitoring is assumed to be accurate but costly.

Let $c(\theta, q)$ denote the total cost to the firm of producing the output level q. The regulator's uncertainty about the firm's costs is summarized by the single, real-valued random variable θ, whose realized value is known to the firm when it chooses q but never becomes known to the regulator. For simplicity, we suppose that θ takes on an arbitrarily large but finite number of values indexed by $k = 1 \dots K$ with $\theta_{k_1} > \theta_{k_2}$ if $k_1 > k_2$. A firm for which $\theta = \theta_k$ will be referred to as a "type k" firm.[13]

In the context of pollution control, let q denote the *reduction* in the firm's emissions from a base level e_0 which would be emitted in the absence of regulation.[14] In this case $c(\theta, q)$ represents the minimum cost to the firm of reducing its emissions by the amount q. This cost may include the cost of operating emission control devices, but it can also include the loss in profit incurred when the firm reduces its emissions by optimally adjusting its mix of inputs and outputs.

It is standard[15] to assume that, for all θ_k and all $q > 0$, $c_q > 0$, $c_{qq} \geq 0$, $c_\theta > 0$, $c_{q\theta} > 0$, $c_{qq\theta} \geq 0$, and $c_{q\theta\theta} \geq 0$, where the subscripts denote partial derivatives. (When $q = 0$, these partial derivatives are also assumed to be nonnegative.) Thus, the firm's marginal cost curve is constant or upward sloping for all values of θ, and both total and marginal costs are increasing functions of θ. For simplicity, we also assume that $c(\theta_k, 0) = 0$ for all k.

Let π_k be the (strictly positive) probability that the regulator assigns to the event that $\theta = \theta_k$. As with similar models, the model in this paper can also be interpreted as one where several firms are regulated and π_k denotes the fraction of the firms for which $\theta = \theta_k$. Under this interpretation, the regulator knows π_k, but the type of a specific firm is either unobservable by the regulator or not verifiable so that different types of firm cannot be regulated separately.

The regulatory regime under which the firm operates consists of three elements. First, the regulator sets a possibly nonlinear payment schedule, $T(\hat{q})$, where $\hat{q}$ is the output level which the firm *reports* that it will produce. Throughout the paper, positive values indicate payments from the firm (e.g., taxes) and negative values indicate payments to the firm. The payment $T(\hat{q})$ is an initial "up-front" tax (or subsidy) that does not depend on the firm's actual output level.

For each reported output level, the regulator also commits to monitoring the firm with probability $\lambda(\hat{q})$. If $\lambda(\hat{q}) < 1$, this implies that the regulator can commit to monitoring the firm with a probability no greater than $\lambda(\hat{q})$. This assumption is relaxed in Section 4. Throughout the paper it is assumed that the regulator can commit to monitoring the firm with a probability no less than $\lambda(\hat{q})$.[16]

[13] Using arguments very similar to those in this paper, the optimal regulations for a single firm with a continuum of possible types can be derived. The resulting expressions for the optimal regulations closely resemble the corresponding expressions for a discrete number of types in Sections 3 and 4.

[14] When q is interpreted as a reduction in emissions from a base level e_0, it is natural to impose the constraint $q \leq e_0$. In the rest of the paper, it is assumed, for simplicity, that this constraint is not binding.

[15] See Besanko and Sappington [9].

[16] Since monitoring in this model serves only to deter "cheating" by firms, a regulator would prefer not to carry out his commitment to monitor once cheating has been deterred. A common response is that a regulator will wish to honor present commitments in order to preserve a "reputation" for the future. However, a formal analysis of reputation effects would require a dynamic model that is beyond the scope of the current paper.

The firm's actual output level, q, is not observed by the regulator unless the firm is monitored, in which case the regulator is assumed to measure q with perfect accuracy. When the firm is monitored, it incurs an additional payment, $R(\hat{q}, q)$, that may depend on both the actual and the reported output levels of the firm.

For the purpose of analysis, it is convenient to suppress the reported output levels, which serve only to specify the regulatory regime under which the firm chooses to operate. An equivalent (and standard) reinterpretation of the previously described regulatory regime is that, by reporting an output level, the firm selects an item from a menu of regulatory schemes proposed by the regulator. An item j in the menu is a triplet consisting of: (1) an up-front payment, T_j, (2) a probability of being monitored, λ_j, and (3) a function, $R_j(q)$, that specifies the "ex post" payment incurred by the firm if it is monitored and the actual output is discovered to be q. If there are K types of firm, then, as observed in footnote 18, it will never be optimal for the regulator to offer a menu with more than K items.

In most of the paper, a triplet T_j, λ_j, and $R_j(q)$ constitutes a valid regulatory scheme if it satisfies the restrictions that $0 \leq \lambda_j \leq 1$ and

$$-\overline{R} \leq R_j(q) \leq \overline{T}, \tag{1}$$

where $\overline{R}$ and $\overline{T}$ are exogenous, nonnegative parameters. (Section 4 considers a further restriction on the set of valid regulatory schemes.) $\overline{T}$ is the largest tax or "fine" that a regulator can impose on a firm that is monitored and found to be producing q. Similarly, $\overline{R}$ is the largest subsidy or rebate that the regulator can provide to a firm that produces q. In the absence of such constraints on the ex post payments, the regulator could prevent a firm from deviating from its reported output level at an arbitrarily small cost by using an arbitrarily large fine to punish noncompliance.

The process of regulation is modeled as a three-stage game. In the first stage, the regulator chooses a menu of regulations. In the second stage, the firm simultaneously chooses both an item from the regulatory menu and the actual output level which it will produce. Negative output levels are not permitted. In the third stage, the firm is monitored with a probability that is determined by the firm's choice in the second stage, and, if the firm is monitored, the firm incurs a payment that also depends on its choice in the second stage. The optimal menu of regulations and the optimal output levels for each type of firm are determined by a perfect Bayesian equilibrium of this game.

A type k firm is assumed to be risk neutral and to choose the regulatory scheme k and output level q_k in the second stage in order to minimize its expected costs:

$$T_k + \lambda_k R_k(q_k) + c(\theta_k, q_k).$$

The firm is also allowed to choose "quit" in the second stage, in which case the firm incurs a reservation cost, $\bar{c}$. For simplicity, it is assumed that—because of unmodeled social costs associated with quitting—the regulator will never wish to propose a menu that will induce any type of firm to choose "quit" in stage 2.

Let $B(q)$ denote the dollar value of the social benefits obtained when the output q is produced. For all $q > 0$, suppose that $B(q) > 0$, $B'(q) > 0$, and $B''(q) \geq 0$, where a prime denotes the derivative. Suppose also that $B(0) \geq 0$, $B'(0) \geq 0$, and $B''(0) \geq 0$.

The optimal menu for the regulator to propose in the first stage of the regulation game and the equilibrium choices of each type of firm in the second stage can be determined by solving the following equivalent problem where the regulator chooses both a regulatory scheme *and* an output level for each type of firm in order to maximize expected net benefits subject to the constraints imposed by the optimizing behavior of the firm in the second stage. Suppose that the regulator chooses K regulatory schemes, $\{T_k, \lambda_k, R_k(q)\}$, and K output levels, $q_k \geq 0$, in order to

$$\max \sum_{k=1}^{K} \pi_k \left[B(q_k) - c(\theta_k, q_k) + \gamma(T_k + \lambda_k R_k(q_k)) - m\lambda_k \right] \tag{2a}$$

subject to the following constraints which must be satisfied for all $1 \leq k,\ j \leq K$ and all $q \geq 0$.

$$-\bar{R} \leq R_k(q) \leq \bar{T} \tag{2b}$$

$$0 \leq \lambda_k \leq 1 \tag{2c}$$

$$T_k + \lambda_k R_k(q_k) + c(\theta_k, q_k) \leq \bar{c} \tag{2d}$$

$$T_k + \lambda_k R_k(q_k) + c(\theta_k, q_k) \leq T_j + \lambda_j R_j(q) + c(\theta_k, q), \tag{2e}$$

where k indexes the regulatory scheme and output level that is chosen in equilibrium by a type k firm,[17] and γ and m are positive parameters whose interpretation is discussed in more detail shortly. Cases where the optimal menu contains fewer than K items or where several types of firm find it optimal to choose the same output are handled formally by allowing the output levels and some or all of the elements of the regulatory schemes to be the same for different values of k.[18]

In terms of the original description of the regulatory regime, choosing regulatory scheme k corresponds to reporting output $\hat{q} = q_k$. For all k, $T(q_k) = T_k$ and $R(q_k, q) = R_k(q)$.

[17]It is standard to assume that a type k firm will make the choice specified for it by the optimization in Eqs. (2a) through (2e) even if it is indifferent between that choice and other available choices. With this assumption, the regulator can implement a desired Bayesian equilibrium by proposing in stage 1 the menu of K regulatory schemes that results from the optimization in Eqs. (2a) through (2e).

[18]There is no loss of generality in restricting the regulator to menus with only K items and assuming that each type of firm chooses a single regulatory scheme and corresponding output level with certainty.

Consider an arbitrary proposed equilibrium. For each k, select from among the choices of a type k firm a regulatory scheme and corresponding output level that maximizes the term in square brackets in Eq. (2a). Now assume that each type of firm chooses with probability one the regulatory scheme and output selected for it and delete from the original menu any items which are no longer chosen by any type of firm. Since deleting a regulatory scheme and output level does not affect the feasibility of the remaining regulation-output combinations with respect to the inequalities in Eqs. (2b) through (2e), the set of K remaining regulatory schemes and corresponding output levels still represents a possible set of equilibrium choices that would be made by each type of firm in stage 2. Moreover, these choices produce a level of expected net benefits at least as large as the original menu. The optimization in Eqs. (2a) through (2e) maximizes over all such sets; hence, it selects an equilibrium where the regulator's menu contains at most K items, each type of firm chooses a single menu item and output level with certainty, and the level of expected net benefits is at least as great as in the originally proposed equilibrium.

The parameter m denotes the cost of monitoring the firm's output. For simplicity, this cost is assumed to be a positive number that is independent of the firm's type or chosen output level. Hence, the term $m\lambda_k$ in Eq. (2a) represents the expected cost of monitoring conditional on regulatory scheme k being chosen by the firm.

The term $T_k + \lambda_k R_k(q_k)$ in Eq. (2a) is the expected payment made by the firm under regulatory scheme k. γ is a positive parameter that denotes the social value of a dollar transferred from the firm to the regulator. A common interpretation is that γ represents the social cost of raising a dollar of revenue via taxes.[19] γ can also be interpreted as the marginal social value produced by the regulator's expenditure of a dollar obtained from the firm.[20] The term $\gamma(T_k + \gamma_k R_k(q_k))$ represents the conditional expected social benefit of the transfers made under regulatory scheme k.

As previously noted, inequalities (2b) and (2c) are part of the definition of a valid regulatory scheme. Inequalities (2d) and (2e) result from the optimizing behavior of the firm in the second stage of the game. The inequalities in Eq. (2d) are "participation" constraints which ensure that a type k firm will prefer to make its equilibrium choices in the second stage rather than to "quit." The inequalities in Eq. (2e) are "incentive compatibility" constraints which guarantee that each type of firm prefers its equilibrium choice of a regulatory scheme and output level to any other available choice of regulation and output.

An argument similar to those typically associated with the "revelation principle" shows that, given the structural assumptions of our model, the regulator can achieve at least as large a value of the objective in Eq. (2a) by using the regulatory framework which we consider as he can by using any of a large class of alternative regulations. For example, consider a Bayesian equilibrium of any regulation game where (i) the equilibrium strategy for the regulator is a pure strategy, (ii) the equilibrium strategy for each type of firm is either pure or mixed, and (iii) the combination of a pure strategy for the regulator and a pure strategy for a type k firm specifies: (1) an output level, q_k, for the type k firm, (2) a probability, λ_k, that the type k firm is monitored, (3) an initial payment, T_k, that the type k firm makes before it is monitored, and (4) a payment schedule, $R_k(q)$, that specifies an additional payment which the type k firm makes (or receives) if it is monitored and found to be producing q. Depending on the regulation game, these quantities may be determined by one or several stages of play. Moreover, the game may impose constraints on the allowable combinations of payments, outputs, etc. in addition to the ones considered in this paper.[21]

Other than exogenous parameters and functions, the four items in the previous paragraph are the only factors that enter into the payoff functions of either the regulator or the firm. Moreover, as long as the regulator is unable to observe the firm's costs and the firm is free to choose any output, all the combinations

[19]Caillaud *et al.* [14] discuss various possibilities for the regulator's objective function, including the one adopted in this paper.

[20]As discussed, for example, in Hahn [24] and OECD [35], some European pollution control agencies may use the funds obtained from regulation to subsidize additional efforts to clean up pollution.

[21]An analysis of the possibility that a regulator might wish to offer lotteries over various combinations of initial taxes, output levels, and so on would be somewhat more complicated. See Myerson [34] for a framework that could be adapted to the model in this paper and used to study this question.

$(q_k, \lambda_k, T_k, R_k(q))$ and $(q_j, \lambda_j, T_j, R_j(q))$ which are realized with positive probability in an equilibrium of the game must also satisfy the constraints in Eqs. (2b)–(2e) for all k, j, and q for the same reasons that these constraints must be satisfied in our model. But the regulatory framework that we consider gives the regulator the greatest freedom to choose the payoff-relevant quantities $q_k, \lambda_k, T_k, R_k(q)$, given the constraints (2b)–(2e). Hence, the expected payoff that the regulator can achieve with our framework must be at least as great as the payoff attainable in the equilibrium of any other game where these payoff-relevant quantities also satisfy Eqs. (2b)–(2e).

3. OPTIMAL REGULATION

This section discusses the optimal regulations specified by the solution to the optimization in Eqs. (2a) through (2e). We particularly wish to note the emergence of rebates as an optimal regulatory response and to compare the optimal regulations in the current environment with the optimal regulations in the benchmark case where the regulator can costlessly observe output and impose an arbitrary, nonlinear tax based on the firm's actual output.

Before considering the optimal regulations, it is helpful to introduce the following notation. Let $\tau_k = T_k + \lambda_k R_k(q_k)$. The quantity τ_k is the total expected payment incurred in equilibrium by a type k firm.

As a point of comparison, consider a model where the structural assumptions, such as the specification of net benefits and the firm's costs, are as described in Section 2, but where the regulator can costlessly impose an arbitrary nonlinear tax that is a function of the firm's actual output. Let τ_k also denote the tax payment incurred by a type k firm in this case. The optimal tax, $\tau(q)$, for the regulator to impose is determined in this model by choosing q_k and τ_k for $k = 1 \ldots K$ to maximize the objective in Eq. (3a) (with $m = 0$) subject to the constraints in Eqs. (3d) and (3e) and the constraint that $q_k \geq 0$. For all k, $\tau(q_k) = \tau_k$. For brevity, we will refer to this benchmark case as the "standard model." Although the structural assumptions in our paper differ somewhat from those in Baron and Myerson's [5] paper, the standard model is closely related to Baron and Myerson's model.

It is convenient to solve the optimization in Eqs. (2a) through (2e) in two stages. In the first stage, the desired output for each type of firm, q_k, is held fixed. For each fixed set of outputs, the optimal regulations to implement these output choices are calculated by solving the constrained optimization in Eqs. (2a) through (2e) for T_k, λ_k, and $R_k(q)$. In the second stage, the optimal output levels are determined.

The next two pages sketch the derivation of the optimal regulations for implementing any fixed, feasible set of output levels. Following this sketch, expressions for these optimal regulations are reported in Eqs. (4a) and (4b) and Eqs. (5a) through (5e).[22]

The derivation of the optimal regulations is simplified by rewriting the optimization in Eqs. (2a) through (2e) in the following equivalent but simpler form.

[22] An appendix containing a more detailed derivation of the optimal regulations was omitted at the request of the editor. A working paper which contains this appendix is available upon request from the author.

First note that it must be optimal to let $R_k(q) = \bar{T}$ for all k and all $q \neq q_k$. For when $q \neq q_k$, $R_k(q)$ appears only in Eq. (2b) and on the right-hand side of Eq. (2e). Hence, setting $R_k(q) = \bar{T}$ merely weakens the incentive compatibility constraints on the other variables. In the rest of the paper, suppose that $R_k(q)$ has been set equal to $\bar{T}$ for $q \neq q_k$.

Let $R_k = R_k(q_k)$ denote the equilibrium ex post payment incurred by a type k firm. Once the q_k are set, a regulatory scheme can be specified by the triplet $\{\tau_k, \lambda_k, R_k\}$. Consider a regulator who chooses K such triplets and the corresponding output levels $q_k \geq 0$ in order to

$$\max \sum_{k=1}^{K} \pi_k [B(q_k) - c(\theta_k, q_k) + \gamma\tau_k - m\lambda_k] \tag{3a}$$

subject to the following constraints which must be satisfied for all $1 \leq k,\ j \leq K$.

$$-\bar{R} \leq R_k \leq \bar{T} \tag{3b}$$

$$0 \leq \lambda_k \leq 1 \tag{3c}$$

$$\tau_k + c(\theta_k, q_k) \leq \bar{c} \tag{3d}$$

$$\tau_k + c(\theta_k, q_k) \leq \tau_j + c(\theta_k, q_j) \tag{3e}$$

$$\tau_K + c(\theta_K, q_K) \leq \tau_j - \lambda_j R_j + \lambda_j \bar{T}, \tag{3f}$$

where, as before, k indexes the regulatory scheme and output level that is chosen in equilibrium by a type k firm.

Equations (3a) through (3d) correspond to Eqs. (2a) through (2d). Equations (3e) and (3f) are what remain of the incentive compatibility constraints in Eq. (2e) once redundant constraints have been eliminated. The constraints in Eq. (3e) will be referred to as "masquerade" constraints since they guarantee that a type k firm will not wish to masquerade as a type j firm by choosing regulatory scheme j and producing output q_j. The constraints in Eq. (3f) guarantee that a type K firm will never wish to "cheat" by choosing regulatory scheme j and then producing no output.[23]

For a fixed set of outputs, it turns out that the procedure for deriving the optimal τ_k in the current model is the same as the procedure for deriving the optimal nonlinear tax in the standard model.[24] In each case, the participation constraint for the type K firm (Eq. (3d)) and the masquerade constraints with $j = k + 1$ must bind. By solving the resulting system of equations recursively starting with τ_K, one obtains the expressions for the optimal τ_k given in Eqs. (4a) and (5c). Hence, for the same set of output levels, the expected payment incurred

[23] If a type k firm is going to deviate from the regulatory scheme and output level intended for it by the regulator, then it does so optimally in one of two ways. First, it can masquerade as a type j firm. Such a masquerade cannot be detected by the regulator since he cannot observe θ. Alternately, the firm may choose to risk the maximum fine, $\bar{T}$, by deviating in a detectable way. In such a case, it is optimal for the firm to choose some regulatory scheme j and then produce nothing. But the cost saving from a given reduction in output is greatest for the type K firm. This turns out to imply that if masquerades are deterred for all types of firm and cheating by producing nothing is deterred for the type K firm, then all other forms of "cheating" by firms will also be deterred.

[24] Spence's [41] derivation of optimal nonlinear prices is an example of a similar calculation.

by a type k firm in our model and the tax paid by such a firm in the standard model are the same. (However, the optimal output for a type k firm will not generally be the same in these two models, so that the actual values of τ_k will also differ in the two cases.)

With the output levels held fixed and optimal values for the τ_k determined as described in the previous paragraph, the optimal values for R_k and λ_k are given by the solution to K uncoupled problems. For each k, λ_k is minimized subject to the constraints in Eqs. (3b), (3c), and (3f). It is clearly optimal to set $R_k = -\bar{R}$ so that the constraint on λ_k in Eq. (3f) is as loose as possible. The optimal value for λ_k is then determined by reducing λ_k until the inequality in Eq. (3f) becomes an equality.

As in the usual analysis of the standard model, it is convenient to write the optimal regulations listed in Eqs. (5a) through (5e) in terms of the following sum of cost terms. For a given set of outputs q_k, let S_k denote the sum specified by

$$S_k = c(\theta_K, q_K) + \sum_{j=k}^{K-1} \left[c(\theta_j, q_j) - c(\theta_j, q_{j+1}) \right], \tag{4a}$$

where the usual convention that a summation equals 0 when the lower limit exceeds the upper limit is used. By collecting the terms in S_k involving q_j, this sum can be rewritten in the form

$$S_k = c(\theta_k, q_k) + \sum_{j=k+1}^{K} \left[c(\theta_j, q_j) - c(\theta_{j-1}, q_j) \right]. \tag{4b}$$

It is well known from the analysis of the standard model that the incentive compatibility constraints in Eq. (3e) cannot be satisfied unless $q_k \geq q_{k+1}$ for all k. In addition, the requirement that $\lambda_k \leq 1$ imposes a further constraint on the feasible output levels; the right-hand side of Eq. (5d) must be less than or equal to one. Any set of output levels which satisfies these two sets of constraints is feasible and can be implemented by the regulations listed in Eqs. (5a) through (5e).

Optimal Regulations for Implementing the Feasible Output Levels q_k

$$R_k(q) = \bar{T}, \qquad \text{for } q \neq q_k \tag{5a}$$

$$R_k(q_k) = -\bar{R} \tag{5b}$$

$$\tau_k = \bar{c} - S_k \tag{5c}$$

$$\lambda_k = \frac{S_k}{\bar{R} + \bar{T}} \tag{5d}$$

$$T_k = \bar{c} - \frac{\bar{T}}{\bar{R} + \bar{T}} S_k \tag{5e}$$

Although the parameters m and γ do not appear directly in the above expressions for the optimal regulations, they do affect these regulations indirectly. The quantity S_k is a function of the output levels which the regulator wishes to implement, and, as we shall see shortly, γ and m play an important role in determining the optimal output levels.

Equation (4a) implies that S_k is a weakly decreasing function of k whenever $q_k \geq q_{k+1}$ for all k. Unless $q_k = q_{k+1}$, $S_k > S_{k+1}$. Hence, λ_k is also a (weakly) decreasing function of k. Both T_k, the up-front tax paid by a type k firm, and τ_k, the expected payment incurred by a type k firm, are (weakly) increasing functions of k. In terms of output, a firm with lower costs produces at least as much as a higher cost firm. If the firm produces more (as will often be the case), then it is more likely to be monitored. The expected payment of such a firm is lower than that of a higher cost firm for two reasons. The firm pays a lower up-front tax and is more likely to receive the rebate which occurs when the firm is monitored and found to be producing its reported output.

If a firm that reports q_k "cheats" by producing q instead, then it both looses its rebate and pays a fine, thus incurring the expected cost $\lambda_k[R_k(q) - R_k(q_k)]$. For a given probability of monitoring, deterrence is maximized by making the difference between the fine for noncompliance and a rebate for compliance as large as possible. Hence, Eq. (5a) specifies that any observed deviation from the reported output should optimally be punished with the greatest possible fine. Since fines are not imposed in equilibrium, this threat involves no actual cost.

Equation (5b) specifies that the reward (or rebate) for compliance should also be made as large as possible. Since rebates will actually be paid in equilibrium and transfers to the firm are costly, this policy might appear to involve excessive social costs. However, the up-front tax, T_k, can be increased to cover the expected cost of the rebate.

Since the cost parameter θ is not observable by the regulator, a low-cost firm can always produce less output than the regulator desires by masquerading as a higher cost firm. This consideration determines both the expected payment in our model and the optimal tax in the standard model. In each case, the differences in the payments incurred by lower and higher cost firms provide the smallest incentive needed to induce each type of firm to produce the output which the regulator desires. In particular, the total expected payment incurred in equilibrium by a type k firm is lower than the equilibrium payment of a type $k + 1$ firm by the amount $c(\theta_k, q_k) - c(\theta_k, q_{k+1})$, which is just sufficient to induce the type k firm to produce q_k rather than masquerading as a type $k + 1$ firm.

As noted previously, for the same output levels the formula for τ_k in Eq. (5c) also describes the optimal nonlinear tax in the standard model. The relation between the up-front tax in our model and the tax in the standard model can therefore be studied by combining Eqs. (4a), (5c), and (5e) to produce the following equation which holds for all $k < K$:

$$T_{k+1} - T_k = \frac{\bar{T}}{\bar{R} + \bar{T}}(\tau_{k+1} - \tau_k) = \frac{\bar{T}}{\bar{R} + \bar{T}}[c(\theta_k, q_k) - c(\theta_k, q_{k+1})]. \quad (6)$$

Moreover, $T_k \geq \tau_k$ for all k, and, as long as $\bar{R} > 0$, the inequality is strict for k such that $q_k > 0$. Hence, the up-front tax in our model is "flatter"; that is, it declines more slowly as a function of (reported) output than either the total expected payment in our model or, for the same output levels, the corresponding tax in the standard model.

When $\bar{T}$ is small compared to $\bar{R}$, the optimal up-front tax is almost flat. Hence, in the realistic case where the size of fines for noncompliance is severely limited, the up-front tax will be almost the same for all types of firm, and the incentive for lower cost types to produce more output will be provided almost entirely by variations in the expected rebate that different types of firm receive. Since the optimal rebate is the same for all levels of reported output, the expected rebate varies only because the probability that a firm is monitored varies with the firm's reported output.

At the other extreme, if $\bar{R} = 0$, so that rebates are impossible, then $T_k = \tau_k$. In this case, all of the incentive to produce more output is provided by variations in the up-front tax just as in the standard model.

The size of the up-front tax and the probability that a firm is monitored are related, for if the up-front tax decreases with the reported output level, then the firm has an incentive to cheat by reporting a high output (and so reducing its tax bill) and then producing nothing. In order to deter such cheating, it is optimal to monitor with greater probability a firm that reports a higher output. But this also implies that a firm which reports a higher output receives a higher expected rebate for compliance. The requirements of deterrence, which cause the monitoring probability to increase with reported output, also permit the up-front tax to be flatter than would be the case if this tax were providing all of the incentive for a lower cost firm to produce more output.

By substituting the formulas for the optimal regulations in Eqs. (5a) through (5e) into the objective in Eq. (2a) (or Eq. (3a)) and simplifying the terms involving S_k, one obtains an equation which describes the expected net benefits generated by an optimally regulated firm as a function of the output levels produced by each type of firm,[25]

$$\gamma \bar{c} + \sum_{k=1}^{K} \{\pi_k[B(q_k) - (1 + \tilde{\gamma})c(\theta_k, q_k)] - \tilde{\gamma}\Pi_k[c(\theta_k, q_k) - c(\theta_{k-1}, q_k)]\}, \tag{7}$$

where $\tilde{\gamma} = \gamma + m/(\bar{R} + \bar{T})$ and $\Pi_k = \Sigma_{j=1}^{k-1}\pi_j$. (Note that $\Pi_1 = 0$ and let $c(\theta_0, q) = 0$ for all q.)

Let $q_{K+1} = 0$. With this notation, the optimal output levels for each type of firm are obtained by choosing q_k for $k = 1 \dots K$ to maximize the objective function in Eq. (7) subject to the "monotonicity" constraints that $q_k \geq q_{k+1}$ for all k and

[25]The following equations are useful in simplifying the expression that results from substituting Eqs. (5a) through (5e) into Eq (2a):

$$\begin{aligned} \sum_{k=1}^{K} \pi_k S_k &= \sum_{k=1}^{K} \pi_k c(\theta_k, q_k) + \sum_{k=1}^{K} \pi_k \sum_{j=k+1}^{K} [c(\theta_j, q_j) - c(\theta_{j-1}, q_j)] \\ &= \sum_{k=1}^{K} \pi_k c(\theta_k, q_k) + \sum_{k=2}^{K} \Pi_k [c(\theta_k, q_k) - c(\theta_{k-1}, q_k)]. \end{aligned}$$

The first equality involves a straight substitution from Eq. (4b). The second equality is obtained from the first by collecting all the terms in the double sum that involve q_k.

subject also to the "monitoring" constraint that

$$\lambda_1 = \frac{1}{\bar{R} + \bar{T}}\left[c(\theta_1, q_1) + \sum_{k=2}^{K}\left[c(\theta_k, q_k) - c(\theta_{k-1}, q_k)\right]\right] \leq 1. \tag{8}$$

Since the monotonicity constraints ensure that S_k is a weakly decreasing function of k, the inequality in Eq. (8) is sufficient to guarantee that $\lambda_k \leq 1$ for all k.

The assumption that $c_{qq\theta} \geq 0$ implies that the difference $c(\theta_k, q_k) - c(\theta_{k-1}, q_k)$ is a convex function of q_k for all $k \geq 2$. Hence, the optimization that determines the optimal output levels is a standard problem where a concave objective function is maximized subject to a convex constraint set. We assume in the rest of the paper that a solution to this optimization exists and is characterized by the usual Kuhn–Tucker conditions.

For $1 < k \leq K$, the first-order conditions for the solution to the above problem are given by

$$B'(q_k) = (1 + \tilde{\gamma})c_q(\theta_k, q_k) + \left(\tilde{\gamma}\frac{\Pi_k}{\pi_k} + \frac{\mu}{\pi_k}\right)\left[c_q(\theta_k, q_k) - c_q(\theta_{k-1}, q_k)\right] + \frac{\delta_{k-1} - \delta_k}{\pi_k}, \tag{9a}$$

where $\mu(\bar{R} + \bar{T})$ is the multiplier for the monitoring constraint in Eq. (8) and δ_k is the multiplier for the monotonicity constraint: $q_{k+1} - q_k \leq 0$. For $k = 1$, the first-order condition is given by

$$B'(q_1) = \left(1 + \tilde{\gamma} + \frac{\mu}{\pi_1}\right)c_q(\theta_1, q_1) - \frac{\delta_1}{\pi_1}. \tag{9b}$$

By setting $m = 0$ in Eq. (3a), substituting for τ_k using Eq. (5c), and simplifying, one can observe that, with $\tilde{\gamma}$ replaced by γ, the optimal output levels in the standard model are also determined by maximizing the objective in Eq. (7) subject to the monotonicity constraints. Hence, if the monitoring constraint in Eq. (8) is not binding, then the optimal outputs both in the current model and in the standard model are determined by the same problem except for the value of a single parameter.

In the rest of this section, we focus mainly on the case where $\bar{R} + \bar{T}$ is sufficiently large that the monitoring constraint does not bind. Note, however, that the first-order conditions imply that, when the monitoring constraint does bind, the result is a reduction in the optimal outputs that resembles the output distortions discussed below. Moreover, if the budget of the agency charged with monitoring and enforcement is sufficiently small, then the right-hand side of the inequality in Eq. (8) should plausibly be reduced from 1 to some lower level $\bar{\lambda}$, which makes it more likely that the monitoring constraint will bind. Thus constraints on enforcement budgets may produce large output distortions even in cases where these distortions would otherwise be thought to be small.

When the monitoring constraint is not binding, it is well known that the optimal outputs can often be obtained by solving an unconstrained problem that ignores the monotonicity constraints. For example, if the differences $\theta_k - \theta_{k-1}$ are chosen to be constant and the ratio Π_k/π_k is a nondecreasing function of k (a discrete analog of a standard assumption in the literature on regulation), then the assumption that $c_{q\theta\theta} \geq 0$ is sufficient to ensure that nonnegative output levels which satisfy the first-order conditions for maximizing the unconstrained objective in Eq. (7) also satisfy the monotonicity constraints. Suppose now that the monotonicity constraints are satisfied by the solution to the unconstrained maximization.

A second useful benchmark is the case of a fully informed regulator who observes θ and the firm's output and can set an arbitrary nonlinear tax based on this information. For each k, such a regulator chooses τ_k and q_k to maximize the objective in Eq. (3a) (with $m = 0$) subject only to the participation constraints in Eq. (3d). Assuming an interior solution, the optimal outputs in this case are given by

$$B'(q_k) = (1 + \gamma)c_q(\theta_k, q_k). \tag{10}$$

Since the term in square brackets in Eq. (9a) is positive, the first-order conditions indicate that, for $k > 1$, the optimal outputs in our model will tend to be smaller than the outputs chosen by a fully informed regulator. It is well known that the regulator's inability to observe costs produces an output distortion in the standard model that reduces the output of all but the lowest cost firm. The size of this distortion depends on γ or, in our model, $\tilde{\gamma}$. In addition, since $\tilde{\gamma} > \gamma$, there is a second distortion in our model that reduces even the output of the lowest cost firm.

If the firm is monitored with probability λ, then $m\lambda$ is the expected cost of monitoring and $\lambda(\overline{R} + \overline{T})$ is the expected penalty for cheating. Hence, the ratio $m/(\overline{R} + \overline{T})$ represents the cost to the regulator of imposing one dollar's worth of expected penalty. The social cost of one dollar of transfers, γ, is augmented in our model by the cost of deterrence. Costly monitoring therefore provides an additional rationale for the output distortions predicted by the standard model.

As $\overline{R} + \overline{T}$ becomes large, the monitoring constraint ceases to bind and $\tilde{\gamma}$ approaches γ. If, in addition, γ becomes small, then the optimal outputs produced in both the current and the standard models approach the output levels that would be chosen by a fully informed regulator. However, even when the optimal levels of output are almost the same in these three cases, we have seen that the form of the optimal regulations will be quite different.

4. THE RIGHT TO BE MONITORED

This section considers the case where the regulator's ability to commit to a monitoring policy is limited by the firm's right to be monitored. In particular, a firm that expects a rebate for compliance can request to be monitored.

When the firm has a right to be monitored, firm types will, in general, divide into two categories. Firm types with sufficiently high costs are monitored with a probability less than one and pay an up-front tax that behaves like the tax in the standard model. In particular, all of the incentive for one type of firm in the category to produce more output than another type in the category is provided by variations in the up-front tax. There are no rebates for these types of firm.

The second category consists of firm types with lower costs. A firm in this category receives a rebate and always exercises its right to be monitored. One optimal tax/rebate policy specifies that all types of firm in the second category pay the same up-front tax. Within this category, all of the incentive to produce more output is then provided by variations in the size of the rebate.

Since monitoring is costly, the regulator will prefer to avoid rebates so that the firm does not request to be monitored. However, when the maximum fine, $\bar{T}$, is small, this fine may need to be supplemented by the threat to withhold a rebate in order to deter the firm from cheating. If $\bar{T}$ is sufficiently small, then all types of firm will be offered rebates and, consequently, fall into the second category mentioned above. On the other hand, if $\bar{T}$ is sufficiently large, then rebates will not be needed for deterrence, and all types of firm will be in the first category.

As in the previous sections, regulation is modeled as a three-stage game. In the first stage of the game, the regulator offers a menu of regulatory schemes. The specification of regulatory scheme j is the same as that in Section 2 except that λ_j is replaced by λ_j^{nr}, which denotes the probability that a firm choosing regulatory scheme j is monitored, but only when the firm does not submit a request to be monitored. (nr stands for "no request".)

In the second stage of the regulation game, the firm simultaneously chooses: (1) an item from the regulatory menu, (2) the actual output level that it will produce, and (3) whether or not to submit a request to be monitored. In the third stage, the firm is monitored with probability λ_j^{nr} if it has chosen regulatory scheme j and not submitted a request to be monitored. If the firm has submitted a request to be monitored, then it is monitored with probability one. Except for the firm's new option to submit a request to be monitored, all the assumptions of the model are the same as those in Section 2.

Whether to submit a request to be monitored is a simple choice for the firm. If a firm chooses regulatory scheme j and output level q_i, then it should request to be monitored if and only if $R_j(q_i) < 0$. (Recall that negative quantities denote payments to the firm.)

Let λ_k denote the equilibrium probability that a type k firm is monitored and R_k denote the equilibrium ex post payment incurred by a type k firm. That is, $R_k = R_k(q_k)$, where q_k is the equilibrium output produced by a type k firm. Given the firm's optimal request policy, the presence of a right to be monitored can be modeled as an additional constraint which the firm's policy imposes on R_k and λ_k. The following inequality must be satisfied for all $1 \le k \le K$:

$$R_k(1 - \lambda_k) \ge 0. \tag{11}$$

It is convenient to describe the equilibrium of the above regulation game in terms of an optimization like that in Eqs. (2a) through (2e). In general, the constraints in such an optimization involve both λ_k and λ_k^{nr}. Fortunately, however, we can restrict attention to regulatory schemes for which $\lambda_k = \lambda_k^{\text{nr}}$.[26] Hence, there

[26] If $R_k \ge 0$, then a type k firm will never submit a request to be monitored, and $\lambda_k = \lambda_k^{\text{nr}}$. If $R_k < 0$, then a type k firm will always submit a request to be monitored unless it intends to deviate from its reported output. In this case, it is also optimal for the regulator to set $\lambda_k^{\text{nr}} = 1$ since this probability serves only to deter cheating and does not affect the equilibrium level of net benefits.

is no ambiguity in using λ_k to denote both the equilibrium probability that a type k firm is monitored and λ_k^{nr}.

With this notation, the optimal regulations and output levels are described by an optimization where the regulator chooses K combinations, $\{T_k, \lambda_k, R_k(q)\}$, and K nonnegative output levels, q_k, in order to maximize the objective in Eq. (2a) subject to the constraints in Eqs. (2b) through (2e) and the new constraint in Eq. (11). As before, k indexes the regulatory scheme and output level chosen in equilibrium by a type k firm.

In the rest of this section, the optimal regulations for implementing any given, feasible set of output levels are discussed. Many of these regulations are unaffected by the presence of the constraint in Eq. (11). In particular, it is still optimal to let $R_k(q) = \overline{T}$ when $q \neq q_k$, and the optimal values for τ_k are still described by Eqs. (4a), (4b), and (5c) in the current circumstances. Only the expressions for the optimal values of λ_k and R_k need to be changed.

The derivation of the optimal values for λ_k and R_k closely parallels the derivation sketched in Section 3. With the output levels held fixed and the optimal values for τ_k specified by Eqs. (5c), the optimal values for λ_k and R_k are obtained by solving K uncoupled problems. In each problem, λ_k is minimized subject to the inequality constraint on R_k in Eq. (1), the inequality constraint $0 \leq \lambda_k \leq 1$, the inequality in Eq. (11), and the inequality in Eq. (3f), which can be written in the form

$$S_k \leq \lambda_k(\overline{T} - R_k), \tag{12}$$

where, as before, S_k denotes the sum of cost terms in Eqs. (4a) and (4b). Recall that the S_k are functions of the output levels which are to be implemented. The only change from the derivation in Section 3 is the presence of the additional constraint in Eq. (11).

When $\overline{T} \geq S_k$, it is straightforward to solve for the optimal values of λ_k and R_k. It is clearly optimal to make the difference on the right-hand side of Eq. (12) as large as possible by setting $R_k = 0$. The optimal value for λ_k is then obtained by reducing λ_k until the inequality in Eq. (12) becomes an equality. The optimal monitoring probability obtained in this way, $\lambda_k = S_k/\overline{T}$, satisfies the constraint that $\lambda_k \leq 1$. Note that if $\overline{T} \geq S_k$ for all k, then the resulting regulations are the same as those which occur when a firm does not have a right to be monitored but $\overline{R} = 0$ so that rebates are impossible.

An interesting (and plausible) case occurs when $\overline{T}$ is "small" but $\overline{R} + \overline{T}$ is "large" so that $\overline{R} + \overline{T} \geq S_k > \overline{T}$. Since S_k is a weakly decreasing function of k, this situation is more likely to occur for a lower cost firm.[27]

Since $S_k > \overline{T}$, the constraint in Eq. (12) can only be satisfied if $\lambda_k = 1$ so that R_k can be negative. With $\lambda_k = 1$, the inequality in Eq. (12) and the inequalities in Eq. (1) can be combined in the following equation, which must be satisfied by the optimal rebate for a type k firm:

$$-\overline{R} \leq R_k \leq \overline{T} - S_k. \tag{13a}$$

[27]If $S_k > \overline{R} + \overline{T}$, then there are no feasible values of λ_k and R_k that satisfy Eq. (12). Hence, a feasible set of output values must satisfy the inequality $S_1 \leq \overline{R} + \overline{T}$, which is just the monitoring constraint from Eq. (8) of Section 3.

The assumptions about $\bar{T}$ and $\bar{R} + \bar{T}$ imply that $-\bar{R} \leq \bar{T} - S_k < 0$, so that the optimal policy involves a rebate, and feasible rebates exist.

Combining the optimal formula for τ_k in Eq. (5c) with the accounting identity $T_k = \tau_k - R_k$ produces the following equation for the optimal up-front tax:

$$T_k = \bar{c} - S_k - R_k. \tag{13b}$$

Any combination $\{T_k, R_k\}$ that satisfies Eqs. (13a) and (13b) is an optimal tax/rebate policy for regulating the type k firm.

One reasonable assumption is that the optimal rebate is kept as small as possible so that the regulator avoids collecting money with the up-front tax merely to return it later. In this case, the right-hand inequality in Eq. (13a) becomes an equality and the optimal up-front tax is a constant: $T_k = \bar{c} - \bar{T}$. Every type of firm for which $S_k > \bar{T}$ pays the same up-front tax. Within this category of types, all of the incentive that a lower cost type of firm requires to produce more output is provided by variations in the rebate offered to different types of firm.

As in Section 3, the optimal output levels are obtained by substituting the optimal regulations for each fixed set of outputs into Eq. (2a) and maximizing the resulting objective function subject to the monotonicity constraints that $q_k \geq q_{k+1}$ for all k (letting $q_{K+1} = 0$) and the monitoring constraint in Eq. (8). (See also footnote 27.)

A detailed analysis of the optimal outputs is not necessary to see that if $\bar{T}$ is sufficiently small, for example, $\bar{T} = 0$, then any firm type that produces positive output will be in a category where $S_k > \bar{T}$. On the other hand, if $\bar{T}$ is sufficiently large, then all firms will be in a category where $S_k \leq \bar{T}$.

5. CONCLUSION

This paper develops a model of pollution control which recognizes that regulators are likely to be poorly informed about a specific firm's cost of pollution control, that monitoring a firm's emission of pollutants is costly, and penalties for noncompliance are often limited. We demonstrate that in such a regulatory environment, there are many circumstances where the optimal regulations resemble deposit–refund systems. The paper therefore reinforces and extends the message of Bohm [10], Costanza and Perrings [17], Solow [40], and others that such systems deserve consideration as a tool for dealing with environmental problems.

A number of factors not included in this paper may further enhance the desirability of regulations with deposit–refund features. For example, a firm may be able to monitor its own emissions (in a verifiable way) more cheaply than an outside regulator can. A deposit–refund system would encourage firms to engage in such activity. Moreover, as was noted in passing in Section 3, the amount of monitoring that an enforcement agency can do may be constrained by the agency's budget. A deposit–refund system provides a way to work around this budget constraint by encouraging firms to pay for their own monitoring activities.

On the other hand, we have also shown that even broad, qualitative features of the optimal regulations may depend on the details of how enforcement is limited or what the rights of polluters are. This suggests a note of caution. Although regulations with deposit–refund features have an intuitive appeal, the proper design of an effective system need not be trivial.

REFERENCES

1. S. Baiman and J. Demski, Economically optimal performance evaluation and control systems, *J. Accounting Res. Suppl.* **18**, 184–219 (1980).
2. D. Baron, Regulatory strategies under asymmetric information, *in* "Bayesian Models in Economic Theory" (M. Boyer and R. Kihlstrom, Eds.), North-Holland, Amsterdam (1984).
3. D. Baron, Design of regulatory mechanisms and institutions, *in* "Handbook of Industrial Organization" (R. Schmalensee and R. Willig, Eds.), Vol. 2, North-Holland, Amsterdam (1989).
4. D. Baron and D. Besanko, Regulation, asymmetric information, and auditing, *Rand J. Econom.* **15**, 447–470 (1984).
5. D. Baron and R. Myerson, Regulating a monopolist with unknown costs, *Econometrica* **50**, 911–930 (1982).
6. W. Baumol and W. Oates, "The Theory of Environmental Policy," 2nd ed., Cambridge Univ. Press, London/New York (1988).
7. B. Beavis and I. Dobbs, Firm behaviour under regulatory control of stochastic environmental wastes by probabilistic constraints, *J. Environ. Econom. Management* **14**, 112–127 (1987).
8. B. Beavis and M. Walker, Random wastes, imperfect monitoring, and environmental quality standards, *J. Public Econom.* **21**, 377–387 (1983).
9. D. Besanko and D. Sappington, "Designing Regulatory Policy with Limited Information," Harwood Academic, Chur, Switzerland (1987).
10. P. Bohm, "Deposit–Refund Systems: Theory and Applications to Environmental, Conservation, and Consumer Policy," Resources for the Future, Washington, DC (1981).
11. P. Bohm and C. Russell, Comparative analysis of alternative policy instruments, *in* "Handbook of Natural Resource and Energy Economics" (A. Kneese and J. Sweeney, Eds.), Vol. 1, North-Holland, Amsterdam (1985).
12. K. Border and J. Sobel, Samurai accountant: A theory of auditing and plunder, *Rev. Econom. Stud.* **54**, 525–540 (1987).
13. B. Bower, R. Barre, J. Kuhner, and C. Russell, "Incentives in Water Quality Management: France and the Ruhr Area," Resources for the Future, Washington, DC (1981).
14. B. Caillaud, R. Guesnerie, P. Rey, and J. Tirole, Government intervention in production and incentives theory: A review of recent contributions, *Rand J. Econom.* **19**, 1–26 (1988).
15. B. Caillaud, R. Guesnerie, and P. Rey, Noisy observation in adverse selection models, *Rev. Econom. Stud.* **59**, 595–615 (1992).
16. P. Chander and L. Wilde, "A General Characterization of Optimal Income Taxation and Enforcement," Social Science Working Paper 791, California Institute of Technology (1992).
17. R. Costanza and C. Perrings, A flexible assurance bonding system for improved environmental management, *Ecological Econom.* **2**, 57–75 (1990).
18. M. Cropper and W. Oates, Environmental economics: A survey, *J. Econom. Lit.* **30**, 675–740 (1992).
19. P. Dasgupta, P. Hammond, and E. Maskin, On imperfect information and optimal pollution control, *Rev. Econom. Stud.* **47**, 857–860 (1980).
20. J. Demski, D. Sappington, and P. Spiller, Managing supplier switching, *Rand J. Econom.* **18**, 77–97 (1987).
21. P. Downing and W. Watson, The economics of enforcing air pollution controls, *J. Environ. Econom. Management* **1**, 219–236 (1974).
22. R. Dye, Optimal monitoring policies in agencies, *Rand J. Econom.* **17**, 339–350 (1986).
23. J. Evans, Optimal contracts with costly conditional auditing, *J. Accounting Res. Suppl.* **18**, 108–139 (1980).
24. R. Hahn, Economic prescriptions for environmental problems: How the patient followed the doctor's orders, *J. Econom. Perspect.* **3**, 95–114 (1989).
25. J. Harford, Firm behavior under imperfectly enforceable pollution standards and taxes, *J. Environ. Econom. Management* **5**, 26–43 (1978).
26. J. Harford, Self-reporting of pollution and the firm's behavior under imperfectly enforceable regulations, *J. Environ. Econom. Management* **14**, 293–303 (1987).

27. W. Harrington, Enforcement leverage when penalties are restricted, *J. Public Econom.* **37**, 29–53 (1988).
28. S. Linder and M. McBride, Enforcement costs and regulatory reform: The agency and firm response, *J. Environ. Econom. Management* **11**, 327–346 (1984).
29. A. Malik, Markets for pollution control when firms are noncompliant, *J. Environ. Econom. Management* **18**, 97–106 (1990).
30. A. Malik, Self-reporting and the design of policies for regulating stochastic pollution, *J. Environ. Econom. Management* **24**, 241–257 (1993).
31. L. Martin, The optimal magnitude and enforcement of evadable Pigovian charges, *Public Finance* **39**, 347–358 (1984).
32. N. Melumad and D. Mookherjee, Delegation as commitment: The case of income tax audits, *Rand J. Econom.* **20**, 139–163 (1989).
33. D. Mookherjee and I. Png, Optimal auditing, insurance, and redistribution, *Quart. J. Econom.* **103**, 399–415 (1989).
34. R. Myerson, Optimal coordination mechanisms in generalized principal–agent problems, *J. Math. Econom.* **10**, 67–81 (1982).
35. Organization for Economic Cooperation and Development (OECD), "Pollution Charges in Practice," Paris (1980).
36. I. Ortuno-Ortin, Inspections and externalities, unpublished manuscript, Department of Economics, Universidad de Alicante, Alicante, Spain (1992).
37. J. Reinganum and L. Wilde, Income tax compliance in a principal–agent framework, *J. Public Econom.* **26**, 1–18 (1985).
38. C. Russell, Monitoring and enforcement, *in* "Public Policies for Environmental Protection" (P. Portney, Ed.), Resources for the Future, Washington, DC (1990).
39. K. Seegerson, Uncertainty and incentives for nonpoint pollution control, *J. Environ. Econom. Management* **15**, 87–98 (1988).
40. R. Solow, The economist's approach to pollution and its control, *Science* **173**, 498–503 (1971).
41. A. M. Spence, Multi-product quantity-dependent prices and profitability constraints, *Rev. Econom. Stud.* **47**, 821–841 (1980).
42. D. Spulber, Optimal environmental regulation under asymmetric information, *J. Public Econom.* **35**, 163–181 (1988).