

Dynamic Taxes and Quotas with Learning

Christopher Costello

University of California, Santa Barbara

Larry Karp

University of California, Berkeley

October 14, 2002

Abstract

We compare dynamic taxes and quotas in a stationary environment where a regulator and a non-strategic firm have asymmetric information. The regulator is able to learn about the unknown cost parameter either by using a tax or a quota that is slack with positive probability. With a tax, the information asymmetry is resolved in one period. Optimal learning using a quota is less transparent, though we show that this search problem has a simple solution. In particular, it is never optimal for the regulator to learn gradually. In the first period, he either ignores the possibility of learning, or he tries to improve his information. Regardless of the outcome in the first period, he never experiments in subsequent periods. We use this result to assess the informational advantage of taxes compared to quotas under asymmetric information.

Keywords: quotas, asymmetric information, searching

JEL classification numbers D83, L50

1 Introduction

An extensive literature compares taxes and quotas when nonstrategic firms and the regulator have asymmetric information about abatement costs. Following Weitzman (1974), many papers (including Malcomson (1978), Roberts and Spence (1976), Stavins (1996), Watson and Ridker (1984) and Yohe (1977)) compare the policies when damages are caused by the flow of pollution.¹ Since the regulator does not exactly know the abatement cost function, he cannot predict the aggregate emissions resulting from the tax. Thus, a tax-setting regulator chooses the *level* of marginal abatement cost, and views aggregate emissions as random. A quota-setting regulator, on the other hand, chooses aggregate emissions, and views the resulting level of marginal abatement cost as a random variable. This difference between the two policies is central to the welfare ranking in the papers cited above. Here we extend the literature by studying the different ways in which the two policies enable the regulator to learn about the firms' abatement costs; that is, we focus on the *informational differences* between the tax and quota in a dynamic setting.

The papers cited above assume that the optimal quota is binding with probability 1. When quotas rights are traded efficiently (and firms are heterogenous, so that trade occurs), this assumption means that the equilibrium quota price conveys to the regulator the same information about an industry-wide cost parameter as does the aggregate equilibrium response to the tax. In this case, there is no informational difference between the policies.

We assume that emissions trading does not occur – either because it is forbidden or too costly, or because homogenous firms have no incentive to trade. In this case, the regulator learns nothing about costs if the quota is binding with probability 1. Such quotas are obviously less informative than taxes. Given that emissions trading does not occur, the regulator might want to use a quota that is slack with positive probability in order to learn about abatement costs.

Our objective is to compare taxes and quotas when the regulator is able to learn about abatement costs, and is able to change the policy level when he has new information. The basic informational difference between the two policies is straightforward. By observing a positive equilibrium response to a tax, the regulator discovers a point on the firm's marginal abatement

¹More recently, a number of papers, including Hoel and Karp (in press), Hoel and Karp (2000), Newell and Pizer (in press), Karp and Zhang (2001), Karp and Zhang (2002), study the case where damages are caused by the stock rather than the flow of pollution.

cost curve. Under a (non-traded) quota, he discovers a point on the marginal abatement cost curve if and only if the quota is slack. In this case, he learns the point at which abatement costs are minimized. Thus, there is an obvious sense in which the tax is more informative than the quota.

Information about a firm's (or an industry's) current abatement cost schedule is valuable only if costs are correlated over time. If costs are imperfectly correlated, the asymmetry of information may diminish but never disappear. In order to emphasize the informational difference between the policies, we assume that the regulator is imperfectly informed about a single cost parameter which is constant over time. Once the regulator learns this parameter, the asymmetry of information vanishes. This assumption makes the optimal policy under taxes extremely simple. Under quite general circumstances, any tax that induces positive output enables the regulator to learn the unknown cost parameter. Thus, the regulator chooses the tax to maximize expected welfare in the current period only. He knows that in the future he can set the tax at the first best (complete information) level.

The optimal policy under quotas is not transparent. Our major contribution is to show that the optimal quota policy is in fact very simple. Once we know the form of this policy, we can easily compare taxes and quotas, and identify the extent of the informational advantage of the former. The next section discusses the problem under quotas and describes the solution. The subsequent section presents the formal model and an example. We then prove our main results, characterizing the optimal quota policy. We use these results to identify the informational advantage of taxes. In a numerical example, we illustrate the parameter space over which taxes or quotas dominate in this dynamic setting.

The assumption that the asymmetry of information vanishes once the regulator has learned the single unknown cost parameter is a means of emphasizing the informational difference between the two policies. If we take this assumption as a literal description, the problem may appear vacuous: If the regulator can learn the unknown parameter by a single experiment, then he would already know it if he had been using a tax in the past, or if the firm had been unregulated, or equivalently if the regulator had used a non-binding quota.

This objection has no force if we consider the case of a new industry, an industry currently regulated with a binding quota, or one which has recently experienced a technological or input price shock that leads to an unknown change in abatement costs. For example, the industry might have installed a new technology that is associated with unknown costs of abating pollu-

tion. If there are substantial costs of changing the technology, it is reasonable to think of the abatement cost function as being constant for a long period of time.

If the regulator could change policies quickly and costlessly, he would be willing to incur a large flow cost for a short period of time (resulting in a low total cost) in order to learn the unknown parameter. However, if policies can be changed only after a non-negligible period of time, it may be necessary to incur a substantial cost in order to learn the unknown parameter. In practice, policies are changed at non-negligible intervals, so the optimization problem is not trivial.

2 The optimal quota as a search problem

This section describes the optimal dynamic quota-setting problem as a problem of optimal search. We compare this problem to a familiar search problem; despite their apparent similarities, the solution to the quota and the familiar search problems are very different.

We use an infinite horizon model of a representative polluting firm. In each period the firm pollutes at its privately optimal level if this level is less than the quota; otherwise, the quota is binding. In each period, given his current beliefs about the abatement cost function, the regulator chooses the quota in order to balance two conflicting objectives. He would like to control the amount of pollution, and simultaneously learn about the true costs so that he can choose a more efficient quota in the future. The quota that would be optimal in a one-period problem (the “myopic quota”), is binding with probability 1 by assumption, so a regulator who uses this quota learns nothing about costs. If he lets the firm produce at its privately optimal level, he learns everything about costs. We discuss these and other assumptions in the next section.

In each period, the regulator can use one of three types of quotas. He can use the myopic quota, in which case he maximizes welfare in the current period, but learns nothing about costs. He can use an “aggressive” quota (defined as one which is slack with probability 1), in which case welfare in the current period is low, but the regulator learns the true cost parameter. Finally, he can use a cautious quota that is binding with probability strictly between 0 and 1. In this case, the regulator learns something about abatement costs, but expected welfare in the current period is lower than under the myopic quota.

One possibility is for the regulator to proceed cautiously. That is, he might refine his infor-

mation over a number of periods until he eventually learns the value of the unknown parameter, or decides that further experimentation is too costly. Rob (1991) analyzes a model which has many of the characteristics of ours, and he finds that this kind of cautious approach is optimal. In Rob's setting, a social planner is uncertain about the location of the market demand curve.² The social planner gradually increases production capacity – requiring costly investment – until he learns the true market demand. In general, learning takes place for more than a single period.

Optimal learning in our setting is qualitatively different. We show that if there is any learning, it takes place in a single period. In addition, the aggressive quota is never optimal if the regulator discounts the future; learning takes place only by using a single, cautious, quota. In the first period the regulator might decide to use either the myopic quota, in which case he learns nothing with probability 1, or a cautious quota, in which case he learns something but might not discover the true cost parameter. Whatever the outcome is in the first period, the regulator never experiments a second time.

The difference in optimal behavior in our search model and in Rob's is due to the differences in the costs of searching “cautiously” and “aggressively” in the two models. A cautious search in Rob's model involves a slight increase in capacity. The amount of information obtained from such a search is modest, but so is the cost. Under the quota, however, it is necessary to incur a significant cost in order to acquire even a small amount of information. This fact follows from the assumption that the myopic quota is binding with probability 1. An aggressive search in Rob's model has a higher expected cost than a more cautious search. For example, if the social planner builds twice the capacity needed to satisfy the market (in Rob's model), he incurs a greater loss than if he over-builds by a small amount. In contrast, an increase in a quota that is slack has no effect on welfare.

3 The model

Here we list the assumptions of our model and explain their role. We then use an example to illustrate these assumptions and to describe the basic properties of the model.

²Much of Rob's analysis concerns the competitive equilibrium. However, for our purposes, the relevant material is his treatment of the social planner's problem.

3.1 Assumptions

The firm knows its abatement cost function, which is constant over time. There are neither exogenous nor endogenous (e.g. investment-related) cost changes. This assumption reduces the complexity of the problem, enabling us to understand how the regulator learns under quotas. We also assume that the firm does not behave strategically with respect to the regulator. Our model does not apply if strategic behavior is important.³ In many circumstances firms are genuinely non-strategic. In addition, we assume that social damages (external to the firm) are caused by the flow rather than the stock of pollution, and the damage function is constant. These assumptions eliminate two possible sources of dynamics, enabling us to focus on learning.

The regulator has full information about the cost function up to an unknown parameter θ . This assumption reduces the dimension of the problem, and is standard in models of asymmetric information.

We define the myopic quota as the quota that would be optimal if the regulator ignored the possibility of learning; that is, the myopic quota minimizes the expectation of current dead-weight loss. In addition to the previous assumptions, we adopt:

Assumption 1 *The firm's individually optimal level of emissions is monotonically increasing in its cost parameter, θ .*

Assumption 2 *The myopic quota is binding with probability 1.*

Assumption 3 *(a) For quotas that are binding, the current loss in social welfare is a finite, convex function of the quota. (b) If the quota is not binding, an increase in the quota does not affect social welfare. (c) The welfare loss under full information is 0.*

Assumption 4 *(a) Let $\underline{\theta}$ and $\bar{\theta} < \infty$ be the lower and upper bounds of the support. The distribution function of θ is strictly increasing over $[\underline{\theta}, \bar{\theta}]$; (b) the distribution has a finite number of mass points, and (c) There are no mass points in an open neighborhood of $\bar{\theta}$.*

Assumption 1 states that the cost parameter affects the firm's individually optimal behavior. The regulator can calculate the smallest possible value of θ consistent with a particular quota

³Moledina, Polasky, Coggins, and Costello (in press) compare taxes and quotas in the situation where strategic firms face a naive regulator. In this situation, the arrival of new information causes "ratcheting", i.e. the tightening of the policy.

being exactly binding. The assumption also implies that if the quota is not binding, the regulator learns the true cost parameter. Thus, the regulator is certain to learn the cost parameter if he uses a quota that is slack with probability 1.

Assumption 2 implies that under the myopic quota, the regulator never acquires any information. This assumption is consistent with the literature cited above, where there can be no learning under the optimal one-period quota in the absence of trade. If Assumption 2 does not hold, the regulator has a positive probability of obtaining information under the myopic quota. This change eliminates the “fixed cost” of searching, one of the two differences between our and Rob’s search models. With this change, we conjecture that consecutive searches might be optimal.

Assumption 3 is standard. Part (a) implies that the myopic quota is unique. An increase in a non-binding quota has no effect on the firm’s actions and therefore does not change the deadweight loss, as part (b) states. Part (c) is merely a normalization. The deadweight loss of using a first-best (full information) policy is 0.

We explained in Section 2 that Assumptions 2 and 3b account for the qualitative difference in optimal searching behavior in this setting, compared to Rob’s model, Assumption 2 means that the cost of acquiring a small amount of information is non-negligible. Assumption 3b means that the cost of using a high quota levels off as soon as the quota is slack. The example below illustrates these features.

Assumption 4 collects technical assumptions about the distribution: Assumption 4a means that there are no “gaps” in the support of the distribution – i.e. the support is a connected set. This assumption simplifies the description of the regulator’s future information set, as a subsequent footnote points out. Assumption 4b is used to prove Lemma 1 and Assumption 4c is used to prove Proposition 2.

3.2 The example

We illustrate the model using the linear functions shown in Figure 1. By emitting one more unit of pollution, the firm saves the marginal abatement cost. For the linear case, the firm’s marginal benefit (MB) of polluting at rate q equals $(\theta - q)b$. The regulator knows the slope b but not the parameter θ , which in this setting equals the unregulated level of emissions. The

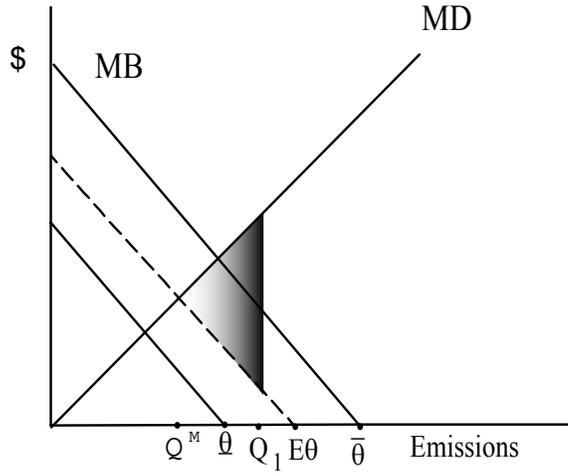


Figure 1: The linear example

expected value of θ is $E\theta$ and the social marginal damage (MD) of pollution is dq .⁴

Under the assumption that the optimal quota is binding with probability 1 (our Assumption 2), Weitzman (1974) shows that it is optimal to use the tax if and only if $b > d$. A simple calculation shows that a necessary condition for the myopic quota to bind with probability 1 is $\frac{b}{d+b}E\theta < \underline{\theta}$.⁵ The expected deadweight loss under the myopic quota equals $\frac{b^2}{2(d+b)}var(\theta)$. By reducing the value of $var(\theta)$, the regulator is able to decrease the expected deadweight loss.

The dashed line in Figure 1 shows the marginal benefit curve associated with $E\theta$. With this linear example, the myopic quota, denoted Q^M , occurs at the intersection of the marginal damage and the expected marginal benefit curve. Consistent with Assumption (2), $Q^M < \underline{\theta}$. Consistent with Assumption (1), the regulator learns the true value of θ if he chooses a quota greater than or equal to $\bar{\theta}$. If the true value of θ happens to be $E\theta$, and if for some reason the regulator uses a quota Q_1 , the deadweight loss in the current period is shown by the shaded

⁴A modification of the example in this section illustrates a situation that violates Assumption 1. Suppose that the horizontal intercept of the firm's marginal benefit curve, θ , is known, but the slope of the curve is private information. In this case, the use of a quota does not enable the regulator to learn about the cost parameter.

⁵Brozovic, Sunding, and Zilberman (2002) use a two part distribution (which violates our Assumption 4a) to show that this condition is not, in general, sufficient to guarantee that the optimal myopic quota binds with probability 1. In the general case, a sufficient condition is $\frac{b\bar{\theta}}{d+b} < \underline{\theta}$. However, our numerical example (section 5) uses a uniform distribution over θ . With this distribution there is a unique local minimum in expected deadweight loss at $\frac{b}{d+b}E\theta$, so $\frac{b}{d+b}E\theta < \underline{\theta}$ is both necessary and sufficient for the myopic quota to bind with probability 1.

area.

It is clear from this figure that the regulator would never choose a quota between Q^M and $\underline{\theta}$. Such a quota increases the current loss in expected welfare without offering any possibility of learning. If the regulator's subjective distribution does not have a mass point at $\underline{\theta}$ he would not want to use a quota close to $\underline{\theta}$, since such a quota provides negligible information but causes a substantial current welfare loss. If he decides to learn about the cost parameter, he uses a quota that is slack with probability strictly greater than 0.

If, for example, the regulator uses $Q_1 > \underline{\theta}$ (Figure 1), then two things can happen. If $\theta < Q_1$ he learns the true value of θ and is able to use the socially optimal quota in the next period; in this case, the future deadweight loss is 0. If $\theta \geq Q_1$ he does not learn the true value of θ , but he learns that the lower bound of the support is higher than he previously thought. He replaces the previous (subjective) lower bound of the support, $\underline{\theta}$, with Q_1 . He updates his subjective probability of θ using Bayes' Rule, improving his ability to use an efficient quota in the next period.

The next section shows that at most one search is optimal and that the aggressive quota is optimal if and only if the regulator does not discount the future. With discounting, the regulator either uses the myopic quota, choosing never to search, and never to learn anything about θ , or he "searches cautiously". In the latter case he uses a quota that is slack with probability strictly between 0 and 1, and in subsequent periods uses the myopic quota conditional on the information obtained during the search. With the cautious search, the regulator improves his information, but might not learn the true value of θ .

4 The general results

In this section we maintain the assumptions described above, but we allow the damage and benefit functions to be general. In view of Assumption 1, we can treat the unknown cost parameter, θ , as the firm's privately optimal level of emissions (in the absence of regulation). This interpretation of the unknown parameter simplifies the proof below, and does not entail any loss in generality beyond Assumption 1. At time t the regulator treats θ as an unknown parameter with support $[\underline{\theta}_t, \bar{\theta}]$.

If $Q_t \leq \underline{\theta}_t$, the regulator's subjective distribution over θ in the subsequent period is unchanged. If $\theta > Q_t > \underline{\theta}_t$ the regulator refines his information, calculating the subjective

distribution in the subsequent period using Bayes' Rule. If the quota binds, $\underline{\theta}_{t+1} = Q_t$.⁶ In every period, the regulator can obtain the current subjective distribution using either the initial distribution or the previous distribution, together with the current value of the lower bound of the support. The lower bound summarizes all of the information that the regulator has acquired since the start of the program; it is the state variable in the regulator's optimal control problem.

We define $L(Q, \theta)$ as the actual deadweight loss in the current period when the regulator uses the quota Q and the true value of the firm's unregulated level of emissions is θ . We define $Q^M(\underline{\theta}_t)$ as the myopic quota (which is unique, by virtue of Assumptions 3a&b):

$$Q^M(\underline{\theta}_t) = \arg \min E_t [L(Q, \theta)].$$

The operator E_t takes the expectation with respect to the unknown parameter θ , given the regulator's beliefs at time t . Assumption 2 implies

$$Q^M(\underline{\theta}_t) < \underline{\theta}_t. \quad (1)$$

The expected value of the deadweight loss in the current period when the regulator uses the *myopic quota*, given the state variable $\underline{\theta}_t$ is $M(\underline{\theta}_t)$:

$$M(\underline{\theta}_t) = E_t [L(Q^M, \theta)].$$

Finally, we define $P(Q, \underline{\theta}_t)$ as the subjective probability that the quota Q is binding, given the value of the state variable, $\underline{\theta}_t$.

By definition, any quota less than $\underline{\theta}_t$ is binding with probability 1. That is

$$P(Q, \underline{\theta}_t) = 1 \quad \forall Q \leq \underline{\theta}_t \quad (2)$$

Assumption (3a) implies

$$E_t L(Q, \theta) > E_t L(Q^M(\underline{\theta}_t), \theta), \quad \forall Q > Q^M(\underline{\theta}_t) \quad (3)$$

Equations (2) and (3) imply that a quota in the interval $(Q^M(\underline{\theta}_t), \underline{\theta}_t]$ increases the deadweight loss in the current period without changing the regulator's information. It is never optimal to use a quota in this interval. If the regulator uses a quota $Q_t > \underline{\theta}_t$ we say that the regulator "searches in period t ". The alternative to searching is to use the myopic quota.

⁶In the absence of Assumption 4a, the updating rule would be a bit more complex. For example, suppose that the density is 0 over the interval (θ_i, θ_{i+1}) and that θ_i and θ_{i+1} are both in the support of θ , with $\theta_i < \theta_{i+1}$. If the quota $Q_t = \theta_i + \varepsilon$ is binding, with $\varepsilon > 0$ and small, then $\underline{\theta}_{t+1} = \theta_{i+1} > Q_t$.

A quota in period t , Q_t , that exceeds $\underline{\theta}_t$ may be slack or binding. In the first case the regulator learns the true value of θ and his problem ends. In the second case he knows that $\theta \geq Q_t$, so $\underline{\theta}_{t+1} = Q_t$. At time t , given the current quota, the regulator knows what the value of the state variable will be in the next period, conditional upon not having learned the true value of θ .

This characteristic holds for an arbitrary number of periods. Given any quota sequence $\{Q_s\}_{s=t}^T$ the regulator knows that by time t' , (where $t < t' \leq T$) either one of the previous quotas will have been slack, or none will have been slack. In the former case he will know the true value of θ . In the latter case, the value of $\underline{\theta}_{t'}$ will be equal to the largest quota between time t and t' . This feature makes our control problem fairly simple to solve.⁷

Providing that the future is discounted, it cannot be optimal to use the myopic quota in one period and then search in a subsequent period. Information does not change under the myopic quota. If it is optimal to incur a cost in order to learn, it is better to do it sooner rather than later.

Rather than looking for an optimal policy function that maps the current state into the current control, we can break the problem into two steps. In the first step, the regulator chooses the number of searches, which we denote T . In the second step, the regulator chooses the optimal conditional quota sequence $\{Q_s\}_{s=t}^T$. This quota sequence is conditional in the following sense: the regulator follows it unless one of the quotas has been slack. If one quota is slack, he learns the true value of θ and he switches to the first best (full information) quota. Hereafter, when we say that a program involves T searches, we mean that the regulator intends to search T more times, conditional upon not learning the true value of θ before the T searches are completed.

Denote the value of the optimal program, i.e. the minimized expectation of the discounted stream of deadweight loss, as $J(\underline{\theta})$. Denote $J^T(\underline{\theta})$ as the value of the optimal program when the regulator decides to search T times. From the previous comments, it is clear that $\min_T J^T(\underline{\theta}) = J(\underline{\theta})$. Our principle result is that it is optimal to search at most one time: the optimal value of T is either 0 or 1.

We begin by showing that the optimal value of T is finite. In order to confirm this fact in

⁷In many stochastic control problems it is difficult to find the optimal (closed loop) control rule because this rule has to specify how the regulator will behave in every possible state at every time in the future. In our problem, the state variable lies in one of two sets. If any of the previous quotas has been slack, the regulator knows the exact value of θ and the problem has ended. In that case, the control rule is trivial: use the full information quota. If none of the previous quotas was binding, the state variable equals the maximum of the previous quotas.

simple manner, we strengthen Assumption 2 slightly, replacing it with⁸

Assumption 5 $Q^M(\underline{\theta}_t) < \underline{\theta}_t - \epsilon$ for some $\epsilon > 0$.

Assumption 2, which implies equation (1), does not exclude the possibility that the myopic quota is arbitrarily close to the lower bound of the support, $\underline{\theta}$. Assumption 5, on the other hand, states that the myopic quota is bounded away from $\underline{\theta}$.

Lemma 1 *Suppose that Assumptions 3, 4a, 4b and 5 hold. Under these assumptions, the optimal T is finite.*

All results are proven in the appendix. The intuition for the lemma is straightforward. Since $\bar{\theta} < \infty$ an infinite number of searches involves an infinite number of marginal searches. The cost of a marginal search is non-negligible, but (provided that there is not a mass point at $\underline{\theta}$) the value of information it reveals is negligible.

We define the regulator's discount factor as β . Lemma 2 considers the case where $\beta = 1$ and Proposition 1 considers the more interesting case where $0 \leq \beta < 1$. If $\beta = 1$ the present discounted value of future deadweight loss is infinite unless the regulator learns the true value of θ . It must therefore be optimal to eventually learn this value. The following lemma states that it is optimal to do so in the first period.

Lemma 2 *If $\beta = 1$ and Assumption 1 and the assumptions of Lemma 1 hold, the optimal value of T is 1 and the optimal quota in the first period is $Q_1 = \bar{\theta}$.*

Any quota $Q_1 \geq \bar{\theta}$ yields the same loss as is therefore also optimal.

Our first main result is:

Proposition 1 *If $0 \leq \beta < 1$ and Assumption 1 and the other assumptions of Lemma 1 hold, the optimal value of T is 0 or 1.*

The proof uses the fact that at time t the regulator can determine the optimal conditional sequence of quotas $\{Q_s\}_{s=t}^T$ for given T . If it were optimal to set $T \geq 2$, then at some point the regulator wants to search two more times. Denote Q_2 as the penultimate search when $T = 2$,

⁸For the linear example, Assumption 5 requires $\frac{b}{d+b}E\theta_t < \underline{\theta}_t - \epsilon$. With a uniform distribution, this inequality implies $\frac{b+2(d+b)\epsilon\bar{\theta}}{2d+b} < \underline{\theta}_t$. Since $\underline{\theta}_t$ is non-decreasing, if this inequality is satisfied at the initial time, it will be satisfied at every subsequent time. Obviously this characteristic holds for many distributions.

and $Q_1 > Q_2$ as the final search. The proof that this sequence cannot be optimal demonstrates that the regulator could have done better by skipping the penultimate search (Q_2) and immediately setting $Q = Q_1$. Here we provide intuition for the result by describing the costs conditional upon certain outcomes; the proof requires taking into account both the probabilities of these outcomes and discounting; it is therefore more complicated than the discussion here.

There are three possibilities: (1) $\theta < Q_2$, (2) $Q_1 < \theta$, or (3) $Q_2 < \theta < Q_1$. In the first case, by Assumption 3b there is no additional loss in the current flow of welfare from using Q_1 instead of Q_2 in the first period; both quotas lead to discovery of the true value of θ . In the second case ($Q_1 < \theta$), the information learned by following the sequence $\{Q_2, Q_1\}$ (i.e. arriving at the conclusion that $\underline{\theta} = Q_1$) could have been obtained at a lower cost by immediately setting Q_1 . Finally, if $Q_2 < \theta < Q_1$, the regulator avoids the cost of the initial search by immediately using Q_1 , and he immediately discovers the true value of θ .

The regulator must choose in the first period one of three types of quotas: the myopic quota (Q^M), the cautious quota ($\underline{\theta} < Q^C < \bar{\theta}$), or the aggressive quota ($Q^A \geq \bar{\theta}$). Our second result shows that the last choice is never used unless $\beta = 1$:

Proposition 2 *Under the assumptions of Proposition 1 and Assumption 4c, the aggressive quota is optimal in the first period if and only if $\beta = 1$.*

Lemma 2 shows that the aggressive quota is optimal if $\beta = 1$. In the case where $\beta < 1$, Proposition 1 guarantees that we need only identify the optimal quota in period 1 (Q_1), since the myopic quota is optimal in all subsequent periods if any uncertainty remains. The proof of Proposition 2 shows that lowering the quota slightly below $\bar{\theta}$ always leads to a decrease in expected deadweight loss. The proof uses the fact that lowering the quota below $\bar{\theta}$ can affect both the expected current and future deadweight loss. A marginal reduction in the quota from $\bar{\theta}$ has no first order effect on either of these terms. It has no second order effect on the future deadweight loss, but it leads to a second order reduction in current deadweight loss. Therefore, a small reduction in the quota (below $\bar{\theta}$) always reduces the present discounted value of expected deadweight loss. Under the assumptions cited in Proposition 2, it is not optimal to resolve all uncertainty except in the limiting case where there is no discounting.

For a fixed positive discount rate, $\beta \rightarrow 1$ as the length of a period during which the policy is fixed approaches 0. If the regulator can change policies frequently, β is close to (but less than) 1. Even in this case it is not optimal to set a quota that is slack with probability 1.

5 The informational advantage of taxes

A tax-setting regulator uses the tax that minimizes expected deadweight loss in the current period, i.e., the myopic tax. He then learns the value of the unknown parameter, and future deadweight loss is zero. Given our characterization of the optimal quota policy, it is straightforward to compare the tax and quota policies and to identify the informational advantage of taxes. We first consider the exogenous characteristics that promote the use of different types of quotas. We then use the linear-quadratic example with additive uncertainty and provide numerical examples.

If the regulator decides to use a quota, a small support for the random variable promotes use of the myopic quota. When the support is small, the amount of uncertainty is small, so the expected deadweight loss of using the myopic quota is small. The cost of a search may nevertheless be large, since it might be necessary to use a large quota (with correspondingly high deadweight loss) in order for there to be a possibility that it is not binding.

If the regulator decides to use a quota and he ignores the future (the discount factor $\beta = 0$) the myopic quota *strictly* dominates a search quota. Since the payoff function is continuous in β (and bounded for all policy choices) at $\beta = 0$, the myopic quota is optimal for small values of β .

For the linear-quadratic example described in Section 3.2, the present discounted expected loss under the myopic quota (J^M) and the optimal tax (J^T) are:

$$J^M = \frac{b^2 \text{var}(\theta)}{2(d+b)(1-\beta)}, \quad J^T = \frac{1}{2} \frac{d^2 \text{var}(\theta)}{d+b}.$$

In the static setting (under Assumption 2), quotas dominate taxes if and only if $b < d$, as Weitzman (1974) showed. In our dynamic setting, the myopic quota dominates taxes if and only if $J^M < J^T$, or $b < d\sqrt{1-\beta}$. Due to their informational advantage, the condition under which taxes dominate the myopic quota is less demanding in a dynamic framework ($d\sqrt{1-\beta} < d$).

Since we do not have a closed form expression for the payoff under the optimal cautious search (quota), we use a numerical example to illustrate the parameter space over which each of the three policies is optimal. The example uses the linear marginal benefit and damage functions from section 3.2, and assumes a uniform distribution for the unknown parameter, θ . We fix the parameters $\bar{\theta} = 24$ and $b = 1$, leaving free the discount factor, β , slope of marginal damages, d , and lower bound on the support of the random variable, $\underline{\theta}$. For four different values

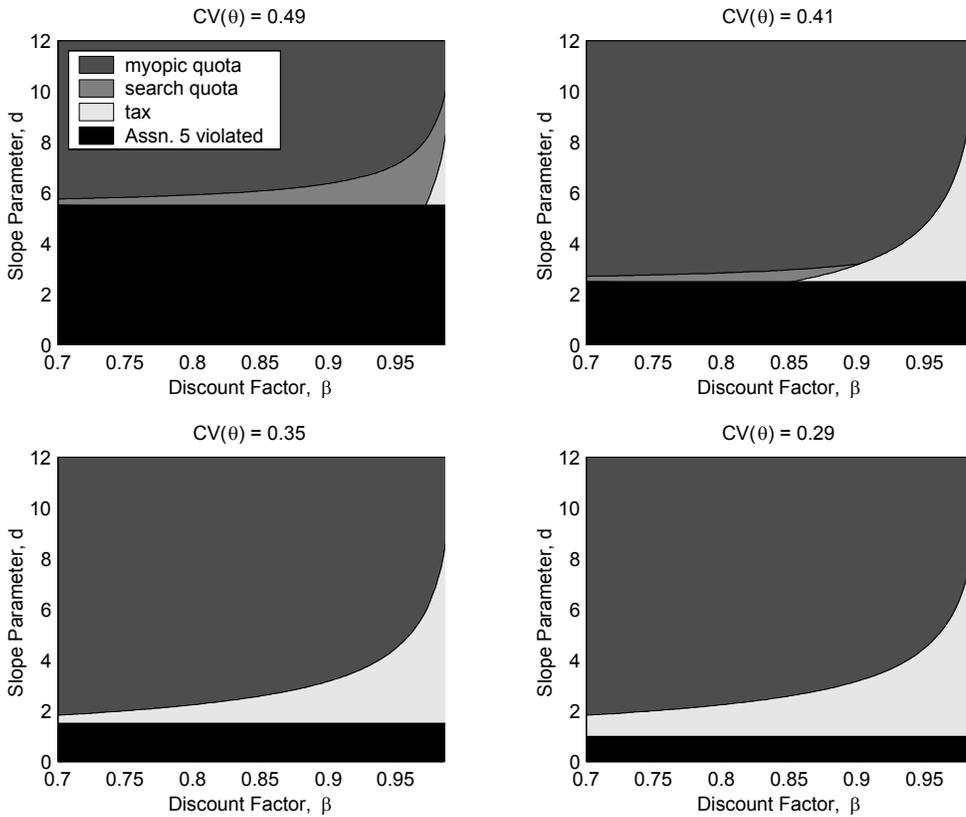


Figure 2: Light shade: tax is optimal. Dark shade: myopic quota is optimal. Middle shade: search quota is optimal. Note: For low d Assumption 5 is violated (black shade).

of $\underline{\theta}$ (2, 4, 6 and 8, corresponding to coefficients of variation equal to 0.49, 0.41, 0.35 and 0.29 respectively) Figure 2 shows the regions of d, β space over which the tax (light shade), myopic quota (dark shade), and search quota (middle shade) are optimal. Sufficiently low values of d violate Assumption 5 and are shaded in black in Figure 2.

We noted above that for low β and high d , taxes are not optimal and for high β and low d , the myopic quota is not optimal. We numerically identify the region of parameter space over which the search quota is preferred over the myopic quota and the tax. Intuition suggests that searching becomes less desirable as uncertainty is reduced. Searching requires a large current investment (a high deadweight loss) in order to reduce future loss. A low level of uncertainty lowers the value of this investment, discouraging the use of the search quota. Figure 2 is consistent with this intuition; for a sufficiently small coefficient of variation, as in the bottom two panels, searching is not optimal for any values of β or d . For a high level of uncertainty,

as in the first panel, a search quota dominates a myopic quota only if d is relatively low. With low values of d , the costs of a search are low. As the discount factor increases, the regulator becomes more willing to incur a current period loss in order to reduce future loss, so the search quota dominates over a wider range of d .

6 Conclusion

A tax has an obvious informational advantage over a quota. The tax can be chosen to minimize current expected deadweight loss, and learning occurs automatically. Under a quota, on the other hand, the regulator may need to sacrifice current expected welfare in order to acquire information and thereby improve future welfare. In order to study the informational differences between the two policies, we need to determine the optimal behavior under a quota.

The quota-setting regulator faces what appears to be a standard search problem, similar to the problem of increasing capacity to discover the size of a market. Two features of the search problem with quotas alter the characteristics of the optimal search. There is a non-negligible cost of acquiring even a small amount of information, and the cost of using a quota that is slack does not increase as the quota increases. These features lead to our two main results. First, gradual search is never optimal. The regulator may or may not search in the first period, but whatever the outcome in that period, he does not search again. Second a quota that leads to discovery of the unknown cost parameter with probability 1 is optimal only in the absence of discounting.

Given the simplicity of the optimal quota program, we can easily rank the tax and quota policies, and determine how the ranking changes from a static to a dynamic context. That is, we can assess the informational advantages of tax. We provided examples using linear functions, and characterized the parameter space over which taxes or quotas dominate.

There are several limitations to our analysis. We have assumed that firms are non-strategic and that there is a single, constant unknown parameter. The problem is fundamentally different if firms are strategic. If there were several unknown parameters, the tax-setting regulator also has a search problem, and the quota-setting regulator has a more complicated search problem. The comparison between the two becomes more complex. Perhaps a more interesting extension is to allow the (single) unknown parameter to be imperfectly correlated over time. The problem for the tax-setting regulator is unchanged, but the search problem for the quota-setting regulator

is much more complex.

7 Appendix: Proofs

Proof. (Lemma 1) Define

$$G(\underline{\theta}_t) \equiv E_t L(\underline{\theta}_t, \theta) - M(\underline{\theta}_t).$$

$G(\underline{\theta}_t)$ is the expected additional cost of using the quota $\underline{\theta}_t$ rather than the myopic quota. Only quotas greater than $\underline{\theta}_t$ have a positive probability of being slack. The regulator learns only by using a quota greater than $\underline{\theta}_t$, so we can regard $G(\underline{\theta}_t)$ as a “fixed cost” of learning. Under assumptions 3 and 5, $G(\underline{\theta}_t) \geq \delta > 0$ for some number δ .

Suppose that at time t , given $\underline{\theta}_t$, the regulator intends to “search” T times (i.e., he intends to use quotas greater than the contemporaneous myopic quota T times, conditional on not having yet learned the value of θ). In this case, the average, over these T searches, of the difference between the quota and the contemporaneous lower bound of the support, is no greater than $\frac{\bar{\theta} - \underline{\theta}_t}{T}$. (If this were not true, then by the time of the T 'th search the quota is greater than the upper bound $\bar{\theta}$.) If $T = \infty$, the regulator plans to use infinitely many quotas that are arbitrarily close to the contemporaneous lower bound. We denote such a quota as a “marginal search”.

In order to show that $T = \infty$ is not optimal, it is sufficient to show that the program associated with $T = \infty$ involves at least one action which is not optimal given the current state. (The existence of such an action violates the Principle of Optimality and therefore cannot be part of an optimal program.)

The current cost of a marginal search is no less than $\delta > 0$. If the subjective distribution does not have a mass point at $\underline{\theta}_t$ the information provided by a marginal search is negligible. In that case, the cost of the search must exceed the value of the information it provides, so such a search cannot be optimal. Since (by assumption) there are a finite number of mass points of the subjective distribution, the program with $T = \infty$ involves infinitely many marginal searches where the cost is strictly positive and the benefit is negligible. ■

Proof. (Lemma 2) By Lemma 1 we need only consider finite values of T . The dynamic programming equation when it is optimal to search one more time is:

$$J^1(\underline{\theta}) = \min_Q [E_1\{L(Q, \theta)\} + K(Q)] \quad (4)$$

where $K(Q)$ is the expected value of future loss given that the current quota equals Q . The expected loss in every future period equals $M(Q)P(Q, \underline{\theta})$. This function is strictly positive if $Q < \bar{\theta}$ and it is equal to zero if $Q \geq \bar{\theta}$. Therefore, in the absence of discounting, we have:

$$K(Q) = \begin{cases} \infty & \text{for } Q < \bar{\theta} \\ 0 & \text{for } Q \geq \bar{\theta}. \end{cases} \quad (5)$$

So in the final period, the optimal choice is to set $Q_1 = \bar{\theta}$, and $J^1(\underline{\theta}) = E_1[L(\bar{\theta}, \theta)]$. Stepping back one period, the DPE when it is optimal to search two more times is:

$$J^2(\underline{\theta}) = \min_Q [E_2\{L(Q, \theta)\} + P(Q, \underline{\theta})J^1(Q)] \quad (6)$$

$$= \min_Q [E_2\{L(Q, \theta)\} + P(Q, \underline{\theta})E_1[L(\bar{\theta}, \theta)]] \quad (7)$$

which is minimized by $Q = \bar{\theta}$. This fact follows because the regulator has committed to learning the true value of θ by the final period.

If, in the penultimate period, he chooses a value $Q_2 < \bar{\theta}$, two things can happen. Either the quota is slack, in which case his loss would have been identical had he chosen $Q_2 = \bar{\theta}$, or the quota is binding. If the quota is binding, he must then choose $\bar{\theta}$ in the final period. But, given that he will acquire the same information by searching either once (by setting $Q_2 = \bar{\theta}$ or twice (by setting $Q_2 < \bar{\theta}$ and $Q_1 = \bar{\theta}$), he can avoid the fixed cost of searching by searching only once. We therefore have a stationary policy function: $Q = \bar{\theta}$ is optimal in every period, and therefore, $Q = \bar{\theta}$ is optimal in the first period. ■

Proof. (Proposition 1) Lemma 2 establishes the result when $\beta = 1$, so we need only consider the case where $0 \leq \beta < 1$. By Lemma 1 we need only consider finite values of T . Suppose, contrary to the Proposition, that the optimal value is $T \geq 2$. In this case, there exists a value of $\underline{\theta}$ at which it is optimal to search exactly two more times (conditional on not first learning the true value of θ). Thus, it is sufficient to show that it cannot be optimal to search two more times.

Suppose to the contrary that for some value of $\underline{\theta}$ it is optimal to search 2 more times. Using previous notation, the dynamic programming equation for this problem is

$$J(\underline{\theta}) = J^2(\underline{\theta}) = \quad (8)$$

$$= \min_Q [E_2\{L(Q, \theta)\} + \beta \{P(Q, \underline{\theta})J^1(Q) + [1 - P(Q, \underline{\theta})]0\}] \quad (9)$$

$$= E_2\{L(Q_2, \theta)\} + \beta \{P(Q_2, \underline{\theta})J^1(Q_2)\}. \quad (10)$$

We use E_i , $i = 1, 2$, to denote the expectation over the unknown parameter θ , conditioned on beliefs at the time when it is optimal to search (at most) i more times; Q_i is the optimal quota when there are (at most) i remaining searches.

If the quota is not binding (which occurs with probability $1 - P(Q, \underline{\theta})$), the regulator learns the true value of θ , and future losses are 0, by Assumption (3c). If the quota is binding, the value of the state in the next period is Q . By the hypothesis that we seek to falsify, it is optimal to search in the next period if the regulator has not learned in the true value of θ . $J^1(Q)$ is the optimal value of the program when it is optimal to search one more time, and the current state is Q .

At the time of the penultimate search, when the regulator uses Q_2 , he knows that θ is either greater or less than Q_2 . Thus, we can write

$$\begin{aligned} E_2 \{L(Q_2, \theta)\} &= [E_{\theta \geq Q_2} \{L(Q_2, \theta)\}] \{P(Q_2, \underline{\theta})\} + \\ &+ [E_{\theta \leq \theta < Q_2} \{L(Q_2, \theta)\}] \{1 - P(Q_2, \underline{\theta})\}. \end{aligned} \quad (11)$$

The expectations on the right side of equation (11) are conditioned on all of the regulator's information at the time of the penultimate search, in addition to the information contained in the inequalities in the subscript of E . Thus, for example, $E_{\theta \geq Q_2}(\cdot)$ is an abbreviation for $E_{\theta \geq Q_2}(E_2(\cdot))$.

Using equation (11) we can rewrite equation (10) as

$$\begin{aligned} J^2(\underline{\theta}) &= [E_{\theta \geq Q_2} \{L(Q_2, \theta)\} + \beta J^1(Q_2)] \{P(Q_2, \underline{\theta})\} + \\ &[E_{\theta \leq \theta < Q_2} \{L(Q_2, \theta)\}] \{1 - P(Q_2, \underline{\theta})\}. \end{aligned} \quad (12)$$

The value function in the next period – the final searching period – (assuming that the regulator has not learned the value of θ) is given by

$$J^1(Q_2) = \min_Q \left[E_1 \{L(Q, \theta)\} + \frac{\beta}{1 - \beta} M(Q) P(Q, Q_2) \right] \quad (13)$$

$$= E_1 \{L(Q_1, \theta)\} + \frac{\beta}{1 - \beta} M(Q_1) P(Q_1, Q_2) \quad (14)$$

Note that the unknown parameter θ appears in only the first term on the right side of equation (14).

Consider an alternative program that involves a single search using Q_1 (rather than Q_2) in the “first period”, when the state is $\underline{\theta}$. Denote the expected value of this program as $A(Q_1, \underline{\theta})$.

We write this expected value by conditioning on the two events: $\theta \geq Q_2$ and $\theta < Q_2$:

$$A(Q_1, \underline{\theta}) = \tag{15}$$

$$= \left[E_{\theta \geq Q_2} \{L(Q_1, \theta)\} + \frac{\beta}{1-\beta} M(Q_1) P(Q_1, Q_2) \right] P(Q_2, \underline{\theta}) + \tag{16}$$

$$[E_{\theta < Q_2} \{L(Q_1, \theta)\}] \{1 - P(Q_2, \underline{\theta})\}$$

Assumption (3b) implies

$$L(Q, \theta) = L(Q', \theta), \quad \forall Q, Q' \geq \theta$$

This equation states that if two quotas are slack, they lead to the same loss in current welfare. The quota used in the last search, Q_1 , involves learning. Therefore $Q_1 > Q_2$. Consequently, we have

$$E_{\theta \leq \theta < Q_2} \{L(Q_2, \theta)\} = E_{\theta \leq \theta < Q_2} \{L(Q_1, \theta)\}.$$

Note that the left and the right side of this equation are, respectively, equal to the last term in square brackets in equation (12) and the last term in square brackets in equation (16). In addition we see (using equation (14)) that the first term in square brackets in equation (16) is equal to $J^1(Q_2)$.

Thus, we can write $A(Q_1, \underline{\theta})$ as

$$A(Q_1, \underline{\theta}) = J^1(Q_2) P(Q_2, \underline{\theta}) + [E_{\theta \leq \theta < Q_2} \{L(Q_2, \theta)\}] \{1 - P(Q_2, \underline{\theta})\} \tag{17}$$

Using equations (12) and (17) we write the difference in payoffs as

$$J^2(\underline{\theta}) - A(Q_1, \underline{\theta}) = [E_{\theta \geq Q_2} \{L(Q_2, \theta)\} + (\beta - 1) J^1(Q_2)] \{P(Q_2, \underline{\theta})\}. \tag{18}$$

If $P(Q_2, \underline{\theta}) = 0$ the regulator learns the true value of θ with probability 1 during the first search. However, by assumption, the regulator intends to search a second time. Therefore $P(Q_2, \underline{\theta}) > 0$. Consequently, to obtain a contradiction we need to show that the term in square brackets in equation (18) is positive. This positive value means that expected discounted stream of costs are higher under the optimal two-search strategy than under an alternative. The two-search strategy is therefore not optimal.

In other words, we must show

$$\frac{E_{\theta \geq Q_2} \{L(Q_2, \theta)\}}{1 - \beta} > J^1(Q_2). \tag{19}$$

Note that $E_{\theta \geq Q_2}(\cdot) = E_{\theta \geq Q_2}(E_2(\cdot)) = E_1(\cdot)$. The first equality is merely a restatement of an earlier abbreviation, and the second equality states that the only additional information that the regulator has at the time of the last search, that he did not have at the time of the penultimate search, is that $\theta \geq Q_2$. Therefore we can rewrite equation (19) as

$$\frac{E_1\{L(Q_2, \theta)\}}{1 - \beta} > J^1(Q_2) \quad (20)$$

The left side of inequality (20) is the present discounted value of using, in perpetuity, the previous quota (which the regulator knows will be binding). However, the costs of that policy exceed the cost of the optimal policy, which requires searching one more time. Consequently inequality (20) is satisfied, and we have a contradiction. ■

Proof. (Proposition 2) Lemma 2 guarantees that the aggressive quota is optimal if $\beta = 1$. Therefore, we need only show that the aggressive quota is not optimal for $0 \leq \beta < 1$.

Define $\Psi(Q)$ as the present value expected loss from setting a quota of Q in the current period (the quantity the regulator wishes to minimize). We decompose $\Psi(Q)$ into the expected loss in the current period, denoted by $\Lambda(Q)$, and the present value of expected future losses, denoted by $H(Q)$: $\Psi(Q) \equiv \Lambda(Q) + H(Q)$. We prove Proposition 2 by showing that $\Psi'(\bar{\theta}) = 0$ and $\Psi''(\bar{\theta}) < 0$; consequently, $\bar{\theta}$ does not minimize $\Psi(Q)$.

Recall that $L(Q, \theta)$ is defined as the actual deadweight loss in the current period when the regulator uses the quota Q and the true value of the firm's unregulated level of emissions is θ . This function is convex by Assumption 3a. We define $\ell(Q, \theta)$ as a differentiable function, convex in Q , with domain $[\underline{\theta}, \theta + \epsilon] \times [\underline{\theta}, \bar{\theta}]$, for $\epsilon > 0$. Using this function we write $L(Q, \theta)$ as

$$L(Q, \theta) \equiv \begin{cases} \ell(Q, \theta) & \text{for } Q \leq \theta, \\ \ell(\theta, \theta) \equiv c(\theta) & \text{for } Q \geq \theta. \end{cases}$$

The deadweight loss when the quota is not binding is $\ell(Q, \theta)$, and the deadweight loss when the quota is binding is $c(\theta)$.

Assumption 2 and the fact that the myopic quota is less than $\bar{\theta}$ implies that $L(Q, \theta)$ is increasing in Q at $Q = \theta - \epsilon$ for small positive ϵ and for θ sufficiently large. (That is, for sufficiently large θ , the full-information optimal quota is binding.) This fact and the definition of $L(Q, \theta)$ implies that $l_Q(\theta, \theta) > 0$ for θ sufficiently large.

The function $L(\cdot)$ is continuous but it is not differentiable at $Q = \theta$; $l(Q, \theta)$, in contrast, is differentiable at $Q = \theta$ by construction. We define $L_Q^-(\theta, \theta)$ as the “left hand” partial derivative

of $L(\cdot)$:

$$L_Q^-(\theta, \theta) \equiv \lim_{\epsilon \rightarrow 0^+} \frac{L(\theta, \theta) - L(\theta - \epsilon, \theta)}{\epsilon} = l_Q(\theta, \theta).$$

The *expected* loss in the current period is:

$$\Lambda(Q) \equiv E_\theta L(Q, \theta) \quad (21)$$

$$= \int_{\underline{\theta}}^{\bar{\theta}} L(Q, \theta) f(\theta) d\theta \quad (22)$$

$$= \int_{\underline{\theta}}^Q c(\theta) f(\theta) d\theta + \int_Q^{\bar{\theta}} \ell(Q, \theta) f(\theta) d\theta \quad (23)$$

Thus

$$\Lambda'(Q) = c(Q) f(Q) - \ell(Q, Q) f(Q) + \int_Q^{\bar{\theta}} \ell_Q(Q, \theta) f(\theta) d\theta$$

$$= \int_Q^{\bar{\theta}} \ell_Q(Q, \theta) f(\theta) d\theta$$

$$\Lambda''(Q) = -\ell_{QQ}(Q, Q) f(Q) + \int_Q^{\bar{\theta}} \ell_{QQ}(Q, \theta) f(\theta) d\theta$$

Evaluating these derivatives at $Q = \bar{\theta}$:

$$\Lambda'(\bar{\theta}) = 0$$

$$\Lambda''(\bar{\theta}) = -\ell_{QQ}(\bar{\theta}, \bar{\theta}) f(\bar{\theta}) < 0.$$

The last inequality uses Assumption 4b, which implies that $f(\bar{\theta}) > 0$. Now we consider the expected future costs of reducing the current quota below $\bar{\theta}$. We define $H(Q)$ as the present value of expected future costs given a quota Q in this period:

$$H(Q) \equiv \frac{\beta}{1 - \beta} P(Q, \underline{\theta}_t) M(Q)$$

This expression equals the probability that the current quota is binding times the discounted stream of payoff in the future (where the lower bound will be $\underline{\theta}_{t+1} = Q$), given that it is optimal to use a myopic quota in the next period (by Proposition 1).

The argument Q in $M(Q)$ determines the amount of uncertainty in the next period, given that the Q was binding. The support of the random variable in the next period is $[Q, \bar{\theta}]$ so as $Q \rightarrow \bar{\theta}$ from below, the amount of uncertainty in the next period vanishes. Consequently,

$$M(\bar{\theta}) = 0 \text{ and } \frac{\partial M(\bar{\theta})}{\partial Q} = 0. \quad (24)$$

The first equality in (24) states that the deadweight loss of using the myopic quota is 0 if there is no uncertainty, i.e. if the myopic and the full information quota are identical. The second equality is implied by the fact that the deadweight loss of using the myopic quota is strictly positive if there is any uncertainty, and it is 0 if there is no uncertainty; thus the deadweight loss is minimized as uncertainty vanishes.

Assumption 4c implies that $P(\bar{\theta}, \cdot) = 0$ and that $P_Q(\bar{\theta}, \bar{\theta})$ and $P_{QQ}(\bar{\theta}, \bar{\theta})$ are finite. Using these facts and equation (24) we have the following derivatives of H :

$$\begin{aligned} H' &= \frac{\beta}{1-\beta}(P_Q M + P M') \\ H'(\bar{\theta}) &= 0 \\ H'' &= \frac{\beta}{1-\beta}(P_{QQ} M + P_Q M' + P_Q M' + P M'') \\ H''(\bar{\theta}) &= 0 \end{aligned}$$

Finally, we have:

$$\begin{aligned} \Psi(Q) &\equiv \Lambda(Q) + H(Q) \\ \Psi'(\bar{\theta}) &= \Lambda'(\bar{\theta}) + H'(\bar{\theta}) = 0 \\ \Psi''(\bar{\theta}) &= \Lambda''(\bar{\theta}) + H''(\bar{\theta}) = -\ell_q(\bar{\theta}, \bar{\theta})f(\bar{\theta}) < 0 \end{aligned}$$

so $\bar{\theta}$ cannot minimize $\Psi(Q)$, and we have our result. ■

References

- BROZOVIC, N., D. SUNDING, AND D. ZILBERMAN (2002): "Prices vs. Quantities Reconsidered," *Giannini Foundation Working paper*.
- HOEL, M., AND L. KARP (2000): "Taxes and Quotas for a Stock Pollutant with Multiplicative Uncertainty," *Journal of Public Economics*, 82, 91 – 114.
- HOEL, M., AND L. KARP (in press): "Taxes Versus Quotas for a Stock Pollutant," *Resource and Energy Economics*.
- KARP, L., AND J. ZHANG (2001): "Controlling a Stock Pollutant with Endogenous Investment and Asymmetric Information," Giannini Foundation of Agricultural Economics, Working Paper No. 928, University of California, Berkeley; <http://are.Berkeley.EDU/karp/>.
- (2002): "Regulation of a stock pollution with correlated cost shocks," Giannini Foundation of Agricultural Economics, University of California, Berkeley; <http://are.Berkeley.EDU/karp/>.
- MALCOMSON, J. (1978): "Prices vs. Quantities: A Critical Note on the Use of Approximations," *Review of Economic Studies*, 45(1), 203–07.
- MOLEDINA, A., S. POLASKY, J. COGGINS, AND C. COSTELLO (in press): "Dynamic Environmental Policy with Strategic Firms: prices vs. quantities," *Journal of Environmental Economics and Management*.
- NEWELL, R., AND W. PIZER (in press): "Stock Externality Regulation Under Uncertainty," *Journal of Environmental Economics and Management*.
- ROB, R. (1991): "Learning and Capacity Expansion Under Demand Uncertainty," *Review of Economic Studies*, 58, 655–675.
- ROBERTS, M., AND A. SPENCE (1976): "Effluent charges and licenses under uncertainty," *Journal of Public Economics*, 5(3-4), 193–208.
- STAVINS, R. (1996): "Correlated uncertainty and policy instrument choice," *Journal of Environmental Economics and Management*, 30(2), 218–232.

- WATSON, W., AND R. RIDKER (1984): "Losses from effluent taxes and quotas under uncertainty," *Journal of Environmental Economics and Management*, 11(4), 310–26.
- WEITZMAN, M. (1974): "Prices versus quantities," *Review of Economic Studies*, 41(4), 477–91.
- YOHE, G. (1977): "Comparisons of price and quantity controls: A survey," *Journal of Comparative Economics*, 1(3), 213–33.