Implementing the Efficient Allocation of Pollution

John Duggan
Department of Political Science
and Department of Economics
University of Rochester
Rochester, NY 14627
U.S.A.

Joanne Roberts
Department of Economics
University of Toronto
Toronto, ON M5S 3G7
Canada

September 23, 1999

Abstract

We provide simple mechanisms to implement the efficient allocation of pollution, first assuming complete information across firms and then allowing for incomplete information. Both mechanisms operate by inducing firms to monitor one another, using firms' reports to determine pollution allotments and transfers. The complete information mechanism determines a firm's transfer according to a linear pollution price, set by other firms. Both mechanisms can be modified to be budget-balanced, in and out of equilibrium. The complete information mechanism implements the efficient allocation even when firms are allowed to use mixed strategies. Under a "single-crossing" condition, the same holds for the incomplete information mechanism.

1 Introduction

Markets to allocate pollution rights play an important role in current efforts to control pollution efficiently — the Environmental Protection Agency has allowed firms to trade pollution permits since 1977 — and this role is likely to grow as pollution abatement efforts intensify. It is clear that a solution to the pollution control problem must involve decentralization of some sort: efficient pollution emissions depend on revenue and cost characteristics of firms, which are typically unknown to regulatory authorities. And with decentralization there arises the possibility of strategic behavior on the part of firms. The extensive literature on the theory of implementation is concerned with precisely this problem, but the mechanisms proposed there, while general in scope, are typically quite complex. In this paper, we construct simple mechanisms to solve the specific problem of efficiently allocating pollution emissions among a fixed set of firms, assuming the regulatory authority can observe pollution emissions and knows the social cost of pollution. The regulator need not know the characteristics of the firms. We initially assume that firms have complete information about each other's characteristics, and we then relax that assumption to allow for incomplete information.

In much of the existing pollution control literature, it is assumed that the regulator knows the efficient level of total pollution — or at least an appropriate "target" level — and the problem is only to allocate that given quantity of pollution permits among firms. As discussed by Dale (1968), a competitive market for pollution permits is well-suited for this task: in equilibrium, firms equate the marginal benefit of pollution to a common price and, therefore, marginal benefits are equated across firms. If the impact of pollution is independent of its source, an efficient allocation of pollution permits is achieved. Kwerel (1977) shows how a combination of licensing and effluent charges can induce firms to reveal their technological characteristics truthfully, allowing the regulator to determine the efficient level of total pollution, but his result relies on price-taking behavior in the market for permits, an untenable assumption if, as Hahn (1984) supposes, some firms have market power. Lewis and Sappington (1995) generalize the simplest problem by allowing firms to have incomplete information but still assume the regulator knows the socially optimal level of pollution.[1]

Our mechanisms endogenously produce the efficient level of total pollution and allocates this total efficiently, while recognizing the strategic incentives of firms to exercise market power, i.e., to influence the prices they pay for pollution emissions. To address this problem, the complete information mechanism replicates the most important feature of competitive markets: each firm purchases a quantity of pollution

---

[1] Xepapdeas (1991), Kritikos (1993), and Herriges, Govidasamy, and Shogren (1994) consider the problem of abatement monitoring in an imperfect information environment with budget–balancing. However, they also assume that the regulator knows the optimal pollution level.

at a price that is independent of the firm's actions. In contrast to the above work, we allow the social impact of pollution to depend on the firm that produces it,[2] so efficiency cannot always be achieved with a uniform price. Nonetheless, by designing incentives for the firms to monitor each other, each firm's price is set appropriately and the allocation of pollution is determined efficiently in the unique pure strategy Nash equilibrium of the mechanism. The incomplete information mechanism still induces the firms to monitor each other, but, because firms may have insufficient information to set prices accurately, pollution emissions are no longer allocated according to a system of linear prices.

The complete information mechanism not only produces the efficient allocation of pollution as an equilibrium outcome, but — since the equilibrium is unique — it ensures that no other allocations can arise as a result of equilibrium behavior. Thus, we implement the efficient allocation in Nash equilibrium using a mechanism that is especially simple compared to those in the implementation literature.[3] Problematic constructions, such as integer games, modulo games, and other forms of "unwinnable competition," are not used. Firms simply select quantities of pollution, used to set prices, and are charged accordingly. The outcome function of the mechanism is continuous, and it is therefore robust to small mistakes in the strategic choices of firms. As long as the regulator can place an upper bound on the efficient level of total pollution (which we will assume), the strategies of firms can also be restricted to compact sets. Duggan and Roberts (1998) show that, under these conditions, the equilibrium outcomes of the mechanism are robust to "small" departures from the complete information assumption. We show that the mechanism is easily adapted to handle mixed strategies and to produce a balanced budget, both in and out of equilibrium. Moreover, we can extend the mechanism to allow for negative externalities across firms as the result of pollution emissions.

Though our analysis takes place in the context of firms and pollution emissions, it applies equally well to the general problem of implementing social welfare optima in quasi-linear environments. Thus, our complete information mechanism is related to the "Nash-efficient" mechanisms surveyed by Groves (1979), and our extension to the case of negative externalities is reminiscent of Hurwicz's (1979) and Walker's (1981) mechanisms for implementing Lindahl equilibria in public good economies. Moore and Repullo (1988) and Varian (1994) propose simple multi-stage mechanisms and, in contrast to other work cited here, use the refinement of subgame perfect equilibrium to implement efficient outcomes.[4] While our approach is distinguished by the specific

---

[2]The assumption of "anonymous" pollution is unrealistic, for example, if firms are geographically distinct and pollution is localized: the social cost of a medium amount of pollution, spread very thinly, may be insignificant; when concentrated at just one locality, however, it may be quite costly.

[3]See Maskin (1977) and Moore and Repullo (1990) for general analyses of Nash implementation.

[4]Eyckmans (1997) adapts Varian's mechanism to implement a proportional solution to a complete

way in which announcements determine prices, our approach also differs from that taken in these papers by explicitly allowing the presence of social externalities (here, the cost of pollution) due to the agents' actions.

The incomplete information mechanism implements the efficient allocation of pollution in Bayesian equilibrium. That this is possible follows from the general analysis of Matsushima (1993), who shows that Bayesian monotonicity, a critical condition for implementation,[5] is automatically satisfied in quasi-linear environments satisfying a condition called "No Consistent Deceptions." This condition restricts the distribution of firms' characteristics and is somewhat weaker than the condition employed in this paper. While Matsushima proves the *possibility* of implementation, he does not provide a simple mechanism suitable for the task. As before, our mechanism gives firms incentives to monitor each other, but, because information is incomplete, one firm can no longer give a point prediction of another's pollution emissions — using a method of belief revelation in Duggan (1998), we ask each firm to report its *beliefs* about its neighbor. The mechanism can be adapted to achieve budget-balance, both in and out of equilibrium, and, if a "single-crossing" property is imposed, efficiency is maintained even if the firms can use mixed strategies.

Dasgupta, Hammond, and Maskin (1980) consider the problem of efficient pollution control, allowing, as we do, for differential impact of pollution and for incomplete information. Their mechanism, a simple adaptation of the mechanism of Groves (1973), Clarke (1971), and Vickrey (1960), has the advantage that firms have dominant strategies leading to efficient pollution. It is well known, however, that the Groves-Clarke-Vickrey (GCV) mechanisms are not generally budget-balanced, and the Dasgupta-Hammond-Maskin mechanism inherits that flaw. Our incomplete information mechanism uses GCV transfers, but only off the equilibrium path. Thus, when there are at least three firms, it is a simple matter to balance the budget in equilibrium. Modifying the mechanism somewhat further, we achieve budget-balance out of equilibrium as well.

Crémer and McLean (1985, 1988) and McAfee and Reny (1992) consider a general mechanism design problem in which agents have incomplete information and offer a mechanism that, under a weak informational condition, generates the optimal allocation as a Bayesian equilibrium outcome and extracts all surplus from the agents. Their restriction on the distribution of agents' types neither implies nor is implied by the condition we impose. Indeed, the objective of these papers is essentially different than ours: they seek to design a mechanism supporting a particular sort of allocation (surplus-extracting) as an equilibrium, but, as noted by Crémer and McLean, their

information pollution abatement problem. This solution requires that individuals bear abatement costs in proportion to their willingness to pay for abatement.

[5]See Jackson (1991) for a general analysis of Bayesian implementation.

mechanism may have multiple equilibria, some of them inefficient. Our objective, in contrast, is to design a mechanism with an efficient *and unique* equilibrium, avoiding the problem of selecting from multiple equilibria. Less critically, we are concerned with achieving budget-balance, while these papers achieve a net flow of transfers, equal to the total surplus of the agents, to the planner.

In Section 2, we describe our complete information model; we present our first mechanism and show that its unique Nash equilibrium yields the socially optimal allocation of pollution; and we discuss possible extensions of our model, some mentioned above. For another example, while we focus on the problem of negative social externalities in this paper, our mechanism works equally well in the "dual" problem of positive social externalities, where firms produce a social *good* as a byproduct of their actions. In Section 3, we describe the incomplete information model and mechanism; we show that it uniquely implements the efficient level of pollution; and we extend our results to account for mixed strategies and budget-balancing.

We close this section by mentioning several issues that, although beyond the scope of this work, merit future consideration. First, while we do not assume that the regulator knows the cost and revenue characteristics of firms, we do assume the regulator observes the pollution outputs of each firm — but in many situations only the aggregate level of pollution may be observed and may not be easily attributed to the firms separately. Thus, we have focused on one important type of asymmetric information in regulatory problems (adverse selection) while abstracting away from another (moral hazard). Second, we have not considered the issue of collusion, which would be a particularly relevant issue when the number of firms is small and firms interact over long time horizons. One approach to this issue would be to design mechanisms that implement the efficient allocation of pollution not only in Nash equilibrium but also in strong Nash equilibrium, which captures the incentives of coalitions to engage in cooperative behavior.[6] Last, as just hinted, pollution control is a *dynamic* problem, and the analysis of efficient pollution control should be extended to cover situations in which output/pollution decisions are made repeatedly over time.

2 Complete Information

2.1 The Model

We consider $n \geq 2$ firms, indexed by i. Denote i's level of pollution (or equivalently i's quantity of pollution permits) by Q_i. The monetary benefit that i receives from producing Q_i units of pollution is denoted $B_i(Q_i)$, and $C(Q_1, \ldots, Q_n)$ is the social cost, measured in monetary terms, imposed on society by the firms' pollution. We assume

---

[6]See Maskin (1979) for a general analysis of implementation in strong Nash equilibrium.

that each $B_i(\cdot)$ is concave and differentiable, that $C(\cdot)$ is continuously differentiable,[7] and that $[\Sigma_{i=1}^n B_i(Q_i)] - C(Q_1, \ldots, Q_n)$ is strictly concave. The benefit and cost functions are common knowledge among the firms, while the regulator is assumed to know only the form of the cost function.

The regulator's problem is to implement the socially optimal allocation of pollution, i.e., the solution to

$$\max_{Q_1,\ldots,Q_n} [\Sigma_{i=1}^n B_i(Q_i)] - C(Q_1, \ldots, Q_n)$$

$$\text{s.t. } Q_1 \geq 0, \ldots, Q_n \geq 0.$$

We assume this problem has a solution, which, by strict concavity, must be unique. Denote it $(Q_1^*, \ldots, Q_n^*)$. We impose the appropriate Inada-type conditions on $C(\cdot)$ and each $B_i(\cdot)$ to ensure an interior solution,[8] so that the social optimum is characterized by the condition that each firm's marginal benefit equals the marginal social cost of pollution:

$$\frac{dB_i}{dQ_i}(Q_i^*) = \frac{\partial C}{\partial Q_i}(Q_1^*, \ldots, Q_n^*)$$

for all i. Lastly, we assume the regulator knows some bound K such that $Q_i^* < K$ for all i.

Because social costs may depend differentially on the emissions of different firms, the efficient allocation cannot generally be obtained by fixing a price common to all firms. Optimality could be achieved if the regulator were to set price $\frac{\partial C}{\partial Q_i}(Q_1^*, \ldots, Q_n^*)$ for each firm i, but this requires knowledge of the social optimum itself. And this in turn requires a familiarity with the firms' benefit functions that is unlikely to be found in practice. In the next subsection, we construct a simple mechanism that implements the efficient allocation of pollution, without presuming such familiarity on the part of the regulator.

2.2 The Mechanism

The mechanism is defined as follows. Firm i purchases a quantity $\hat{Q}_i \in [0, K]$ for itself and reports a quantity $\overline{Q}_{i-1} \in [0, K]$ for its "neighbor," firm $i-1$, where we treat n as firm 1's neighbor. As a function of these reports, firm i pays

$$\hat{Q}_i \frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n) + |\overline{Q}_{i-1} - \hat{Q}_{i-1}|.$$

---

[7]Continuity of the derivative is used only to ensure that the mechanism defined in Section 2.3 is continuous.

[8]Specifically, $\lim_{Q_i \to 0}[\frac{\partial B_i}{\partial Q_i}(Q_i) - \frac{\partial C}{\partial Q_i}(Q_1, \ldots, Q_n)] = \infty$, for all $Q_1, \ldots, Q_n$.

Thus, firm i faces the price $\frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n)$, which is independent of its own reports. The second term in the firm's payment is a penalty for misrepresenting the demand of its neighbor.

Proposition 1. *The unique pure strategy Nash equilibrium of the above mechanism is given by* $(\overline{Q}_{i-1}, \hat{Q}_i) = (Q_{i-1}^*, Q_i^*)$ *for all i, and this yields the socially optimal allocation* $(Q_1^*, \ldots, Q_n^*)$ *of pollution.*

It is straightforward to verify, using concavity of each $B_i(\cdot)$, that the above specification of strategies is indeed a Nash equilibrium. To verify that it is unique, consider an arbitrary pure strategy equilibrium, and note that, because firm i cannot affect its own price, $\hat{Q}_i$ solves

$$\max_{Q_i \in [0,K]} B_i(Q_i) - Q_i \frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n).$$

By our assumptions on $B_i(\cdot)$, $\hat{Q}_i > 0$ and

$$\frac{dB_i}{dQ_i}(\hat{Q}_i) \geq \frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1} \ldots, \hat{Q}_n)$$

are satisfied. Assume for now that the first order condition holds with equality for each i.[9] Also note that $(\overline{Q}_i, \hat{Q}_{i+1})$ is a best response for firm $i+1$ only if $\hat{Q}_i = \overline{Q}_i$, so we have

$$\frac{dB_i}{dQ_i}(\hat{Q}_i) = \frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \hat{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n)$$

for all i, which is satisfied only at the social optimum.

Now suppose that the first order condition of some firm i is not met with equality. Writing

$$W(Q_1, \ldots, Q_n) = \left[\sum_{i=1}^n B_i(Q_i)\right] - C(Q_1, \ldots, Q_n),$$

this is equivalent to $\frac{\partial W}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_n) \neq 0$ and implies that $\frac{\partial W}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_n) > 0$ and $\hat{Q}_i = K > Q_i^*$. Thus, $\nabla W(\hat{Q}_1, \ldots, \hat{Q}_n) > 0$ and

$$\begin{aligned} 0 &> \nabla W(\hat{Q}_1, \ldots, \hat{Q}_n) \cdot [(Q_1^*, \ldots, Q_n^*) - (\hat{Q}_1, \ldots, \hat{Q}_n)] \\ &\geq W(Q_1^*, \ldots, Q_n^*) - W(\hat{Q}_1, \ldots, \hat{Q}_n), \end{aligned}$$

where the weak inequality follows from concavity of $W(\cdot)$ and Sundaram's (1996) Theorem 7.9. But then $W(\hat{Q}_1, \ldots, \hat{Q}_n) > W(Q_1^*, \ldots, Q_n^*)$, a contradiction.

---

[9]If the firms were not restricted to the compact set $[0, K]$, equality would obviously obtain. Dropping that restriction would simplify the proof but would result in a less "well-behaved" mechanism.

2.3 Extensions

The model of Section 2.1 and the result of Section 2.2 can be extended in several interesting ways.

Mixed Strategies. Proposition 1 restricts firms to pure strategies, but the result can be extended if we assume each $B_i(\cdot)$ is *strictly* concave. In this case, consider an arbitrary mixed strategy equilibrium. Then $\frac{\partial C}{\partial Q_i}$ is a random variable, and firm i's strategy can put positive probability only on solutions to

$$\max_{Q_i \in [0,K]} B_i(Q_i) - Q_i E\left[\frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n)\right].$$

Since $B_i(\cdot)$ is strictly concave, this problem has a unique solution, say $\hat{Q}_i$. Thus, i's mixed strategy is to play $\hat{Q}_i$ with probability one. A similar observation holds true for the other firms, so the argument of Section 2.2 applies.

Budget-balancing. When $n \geq 3$, the mechanism is easily modified to achieve budget-balance in equilibrium: simply subtract from firm i's payment to the regulator the amount

$$\hat{Q}_{i+1}\frac{\partial C}{\partial Q_{i+1}}(\hat{Q}_1, \ldots, \hat{Q}_{i-1}, \overline{Q}_i, \overline{Q}_{i+1}, \hat{Q}_{i+2}, \ldots, \hat{Q}_n),$$

where we use $\overline{Q}_i$ here instead of $\hat{Q}_i$ and we drop the term $|\overline{Q}_i - \hat{Q}_i|$, so firm i cannot affect this adjustment. Since $\overline{Q}_i = \hat{Q}_i$ and $\overline{Q}_{i+1} = \hat{Q}_{i+1}$ in equilibrium, this amount equals firm $i + 1$'s equilibrium payment, yielding a balanced budget.

To balance the budget out of equilibrium as well, we modify the original mechanism somewhat. In addition to purchasing quantity $\hat{Q}_i$, firm i reports quantities $\overline{Q}_{i-1}$ and $\overline{\overline{Q}}_{i-2}$ for two neighbors. As a function of these reports, firm i pays the "base amount"

$$\hat{Q}_i\frac{\partial C}{\partial Q_i}(\hat{Q}_1, \ldots, \hat{Q}_{i-2}, \overline{\overline{Q}}_{i-1}, \overline{Q}_i, \hat{Q}_{i+1}, \ldots, \hat{Q}_n) + |\overline{Q}_{i-1} - \hat{Q}_{i-1}| + |\overline{\overline{Q}}_{i-1} - \hat{Q}_{i-2}|.$$

As with the original mechanism, firm i's reports of $\overline{Q}_{i-1}$ and $\overline{\overline{Q}}_{i-2}$ will match $\hat{Q}_{i-1}$ and $\hat{Q}_{i-2}$ in equilibrium. The difference that allows us to fully balance the budget is that now the first part of the base payment is independent of firm $i - 1$'s reports. Budget-balance is achieved by subtracting from firm i's payment the amount

$$\hat{Q}_{i+1}\frac{\partial C}{\partial Q_{i+1}}(\hat{Q}_1, \ldots, \hat{Q}_{i-2}, \overline{\overline{Q}}_i, \overline{Q}_{i+1}, \hat{Q}_{i+1}, \ldots, \hat{Q}_n) + |\overline{Q}_{i+1} - \hat{Q}_{i+1}| + |\overline{\overline{Q}}_{i-1} - \hat{Q}_{i-1}|,$$

which is independent of firm i's reports. The first term above is exactly the first term in firm $i+1$'s base payment, whereas the second and third terms above are the second term in firm $i + 2$'s base payment and the third term in firm $i + 1$'s base payment.

Negative Externalities Across Firms. We have assumed each firm's benefit from polluting is independent of the levels of pollution of other firms. A more general model would allow for externalities: firms may experience either market externalities (as when high levels of pollution by other firms may reflect high levels of production and a competitive output market) or production externalities as pollution levels rise. We now allow for externalities among the firms, using $B_i(Q_1, \ldots, Q_n)$ to denote the benefit of firm i corresponding to pollution quantities $Q_1, \ldots, Q_n$. We assume externalities are negative, i.e., $\frac{\partial B_i}{\partial Q_j} < 0$ for $i \neq j$, and we impose the appropriate Inada-type conditions to guarantee interior solutions. The regulator's problem is defined as before, with unique solution $(Q_1^*, \ldots, Q_n^*)$ given by the first order conditions

$$\sum_{i=1}^{n} \frac{\partial B_i}{\partial Q_j}(Q_1^*, \ldots, Q_n^*) = \frac{\partial C}{\partial Q_j}(Q_1^*, \ldots, Q_n^*),$$

$j = 1, \ldots, n$.

We modify the mechanism as follows. We have each firm i purchase a vector $\hat{Q}^i = (\hat{Q}_1^i, \ldots, \hat{Q}_n^i)$ of pollution quantities, one quantity for each firm,[10] where $\hat{Q}_j^i$ is interpreted as an amount added to firm j's pollution by firm i. Note that Q_i^i is the amount of firm i's pollution purchased by itself. If $\hat{Q}_j^i < 0$, which we allow, then firm j's outputs are decreased by $-\hat{Q}_j^i$. Along with these purchases, firm i reports a vector $\overline{Q}^{i-1} = (\overline{Q}_1^{i-1}, \ldots, \overline{Q}_n^{i-1})$, where $\overline{Q}_j^{i-1}$ represents the increment to firm j's pollution purchased by firm $i - 1$. Firm i is then allocated the total amount,

$$\mathbf{Q}_i = \max\{0, \sum_{j=1}^{n} \hat{Q}_i^j\},$$

of pollution purchased for it, and is charged for its own purchases according to a vector of prices. Once again, it is important that firm i not be able to influence its prices, so for each k we let

$$\overline{\mathbf{Q}}_k = \max\{0, \overline{Q}_k^i + \sum_{j \neq i} \hat{Q}_k^j\},$$

which is the total amount of firm k's pollution with firm i's increment replaced by $\overline{Q}_k^i$, reported by firm $i + 1$. Firm i then faces price

$$2\frac{\partial C}{\partial Q_i}(\overline{\mathbf{Q}}_1, \ldots, \overline{\mathbf{Q}}_n)$$

for its own pollution output, and faces prices

$$-\frac{1}{n-1}\frac{\partial C}{\partial Q_j}(\overline{\mathbf{Q}}_1, \ldots, \overline{\mathbf{Q}}_n),$$

---

[10]If externalities among firms are limited, we can simplify the mechanism by having firm i only purchase pollution quantities for firms imposing externalities on i.

$j = 1, \ldots, n$, for other firms' pollution outputs. Note that the latter prices are negative, so firm i pays for the reduction of other firms' pollution outputs and is compensated for increases in their pollution levels. In addition, i pays the penalty $||\overline{Q}^{i-1} - \hat{Q}^{i-1}||$ for misrepresenting its neighbor's pollution purchases.

The argument that this mechanism implements the socially optimal allocation of pollution is similar to our earlier one. Ignoring corner solutions, every equilibrium pollution allocation must satisfy

$$\frac{\partial B_i}{\partial Q_i}(\mathbf{Q}_1, \ldots, \mathbf{Q}_n) = 2\frac{\partial C}{\partial Q_i}(\overline{\mathbf{Q}}_1, \ldots, \overline{\mathbf{Q}}_n)$$

$$\frac{\partial B_i}{\partial Q_j}(\mathbf{Q}_1, \ldots, \mathbf{Q}_n) = -\frac{1}{n-1}\frac{\partial C}{\partial Q_j}(\overline{\mathbf{Q}}_1, \ldots, \overline{\mathbf{Q}}_n)$$

for all i and all $j \neq i$. Because of the strict incentive for each firm i to report its neighbor's purchases accurately, we also have $\mathbf{Q} = \overline{\mathbf{Q}}$ in equilibrium. Making this substitution and summing up the firms' first order conditions with respect to each Q_j, we see that the first order conditions for the socially optimal allocation are met.

Positive Social Externalities. In Sections 2.1 and 2.2, we analyzed the problem of negative social externalities, of which pollution is a special case, but our mechanism also implements solutions to problems of *positive* social externalities. We now interpret Q_i as an activity of firm i that costs $C_i(Q_i)$, where $C_i(\cdot)$ is differentiable and convex, and the social benefit of activity is $B(Q_1, \ldots, Q_n)$, where $B(\cdot)$ is also differentiable. We now assume that $B(Q_1, \ldots, Q_n) - \sum_{i=1}^{n} C_i(Q_i)$ is strictly concave and that the social optimization problem has an interior solution. The mechanism is unchanged, except that payments from firms become payments to firms. For an example of positive social externalities, suppose that Q_i measures the attractiveness of storefronts in a downtown area, $C_i(Q_i)$ is the cost to firm i of maintaining a storefront of quality Q_i, and $B(Q_1, \ldots, Q_n)$ is the corresponding social benefit. Of course, the agents under consideration need not be firms. They may be workers in a factory, where Q_i denotes i's contribution of effort, $C_i(Q_i)$ a cost of effort, and $B(Q_1, \ldots, Q_n)$ the monetary worth of output, as a function of the vector of efforts expended by workers.

Abatement. We have assumed that firm i may emit pollution in quantities not exceeding $\hat{Q}_i$, the quantity of pollution purchased by i. In some situations, it would be equally reasonable to let firm i produce any amount $\tilde{Q}_i$ of pollution, so long as it cleans up the quantity $\tilde{Q}_i - \hat{Q}_i$ of pollution not paid for. Letting $C_i^*(Q_i)$ denote i's cost of abating Q_i units of pollution, we can extend our results to this case. Define

$$B_i^*(Q_i') = \max_{Q_i \in [Q_i', \infty]} B_i(Q_i) - C_i^*(Q_i - Q_i'),$$

which reflects the monetary benefit to firm i of Q'_i pollution permits. Since pollution in excess of Q'_i is cleaned up, it need not be included in the regulator's calculation of social cost, so the maximization problem of Section 2.1 (with B^*_i substituted for B_i) is still the appropriate one. If $B^*_i(\cdot)$ is differentiable and concave, the remainder of our analysis carries over without change.

3 Incomplete Information

3.1 The Model

Each firm i now has a finite set Θ_i of possible characteristics, or "types," denoted θ_i, θ'_i, etc., representing the profitability of pollution to firm i. Let $\Theta = \times_{i \in N} \Theta_i$ be the set of characteristic profiles $\theta = (\theta_1, \ldots, \theta_n)$, and let p denote the joint distribution of the firms' characteristics. For simplicity, we assume $p(\theta) > 0$ for each profile θ. In this section, we can allow firm i's benefit to depend on the pollution levels of all firms: thus, $B_i(Q_1, \ldots, Q_n, \theta_i)$ denotes the profit of firm i when the pollution outputs of the firms are $Q_1, \ldots, Q_n$ and i's characteristic is θ_i. The social cost function is $C(Q_1, \ldots, Q_n)$, as before. We assume that each firm's characteristic is private information, and that the distribution of characteristics, the benefit functions, and the social cost function are common knowledge among the firms. The regulator is assumed to know the distribution, the benefit functions, and the cost function, but not the characteristics of the firms. We assume that the regulator's problem,

$$\max_{Q_1, \ldots, Q_n} \left[\sum_{i=1}^n B_i(Q_1, \ldots, Q_n, \theta_i) \right] - C(Q_1, \ldots, Q_n)$$

$$\text{s.t. } Q_1 \geq 0, \ldots, Q_n \geq 0,$$

has a unique solution, denoted $Q^*(\theta) = (Q^*_1(\theta), \ldots, Q^*_n(\theta))$, for every profile θ of characteristics. Further, we assume the solution is sensitive to the characteristics of the agents: $\theta_i \neq \theta'_i$ implies $Q^*(\theta_i, \theta_{-i}) \neq Q^*(\theta'_i, \theta_{-i})$ for all θ_{-i}. Neither concavity nor differentiability are required in this section.

Once we allow for incomplete information, we must address the issue of incentive compatibility.[11] That is, along with an allocation of pollution, we must specify a vector of payments from the firms for every profile of characteristics such that truthful reporting of characteristics is a Bayesian equilibrium (see Harsanyi (1967-68)). To state this condition formally, we let $t_i(\theta)$ denote the payment from firm i when characteristic profile θ is realized; we let $t(\theta) = (t_1(\theta), \ldots, t_n(\theta))$ be the profile of payments; and we let $p_{\theta_i}(\theta_{-i})$ denote the probability that the characteristic profile

---

[11]Jackson (1991) shows that incentive compatibility is a necessary condition for Bayesian implementation. It is automatically satisfied in complete information environments with quasi-linear preferences, as in Section 2.

$\theta_{-i} = (\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_n)$ is realized, conditional on θ_i. Then *incentive compatibility* is the requirement that, for all i, all θ_i, and all θ_i',

$$\sum_{\theta_{-i} \in \Theta_{-i}} [B_i(Q^*(\theta), \theta_i) - t_i(\theta)] \, p_{\theta_i}(\theta_{-i})$$

$$\geq \sum_{\theta_{-i} \in \Theta_{-i}} [B_i(Q^*(\theta_i', \theta_{-i}), \theta_i) - t_i(\theta_i', \theta_{-i})] \, p_{\theta_i}(\theta_{-i}).$$

In words, if every other firm reports is characteristics truthfully, no type of firm i has an incentive to misreport its type. D'Aspremont and Gerard-Varet (1979) establish the existence of a function $t^*(\cdot)$ such that (Q^*, t^*) is incentive compatible. Moreover, when the characteristics of the firms are independently distributed, they show that t^* can be constructed to be budget-balanced: $\sum_{i=1}^n t_i^*(\theta) = 0$ for all θ. We will assume only that (Q^*, t^*) is incentive compatible for now, though later we entertain the possibility of balanced transfers.

Our main result on incomplete information makes use of the following technical restriction on the distribution of the firms' characteristics — it should be apparent that it holds generically (i.e., for an open, dense set of distributions on characteristic profiles). Below, let q_{θ_i} denote the distribution of θ_{i-1} conditional on θ_i, i.e., q_{θ_i} is the marginal of p_{θ_i} on firm $i-1$'s characteristic.

($\bullet$) For all i, all θ_i, and all θ_i', if A and B are distinct subsets of Θ_{i-1}, then $q_{\theta_i}(A) \neq q_{\theta_i'}(B)$.

The condition can be broken into two parts. First, for a given characteristic θ_i of firm i, no two distinct sets of $i-1$'s characteristics are realized with exactly the same conditional probability. Second, distinct types of firm i never assign identical probabilities to distinct sets of $i-1$'s types. Though ($\bullet$) holds generically, it is potentially restrictive: the first part of the condition rules out uniformly distributed characteristics; the second part rules out complete information.[12] Later, we impose a "single-crossing" property on the firms' benefit functions that allows us to weaken ($\bullet$).

The mechanism of the next section makes use of a pairwise monitoring scheme similar to that of the previous section. Now, however, we ask firms to report their characteristics directly rather than pollution quantities. Because firm i has incomplete information about $i-1$, we cannot ask firm i to report $i-1$'s type. Instead, we ask firm i to report the *distribution* of $i-1$'s *reported* types. Assuming type θ_i of firm i reveals the actual distribution conditional on θ_i, an implication of condition ($\bullet$) is that any deception by $i-1$ will be revealed to the regulator in this distribution. For example,

---

[12]Complete information models, formulated in Bayesian terms, give firms identical sets of types jointly distributed along diagonal of Θ.

suppose there two firms with sets of characteristics $\Theta_1 = \{\theta_1', \theta_1''\}$ and $\Theta_2 = \{\theta_2', \theta_2''\}$ distributed with the following probabilities, which clearly satisfy condition ($\bullet$).

	θ_2'	θ_2''
θ_1'	$\frac{2}{5}$	$\frac{3}{10}$
θ_1''	$\frac{1}{5}$	$\frac{1}{10}$

If firm 1 reports deceptively, say type θ_1' reporting θ_1'' and visa versa, consider the distributions of 1's reports conditional on firm 2's types. Conditional on type θ_2', firm 2 expects 1 to report θ_1' with probability 1/3 and θ_1'' with probability 2/3. If firm 2 reports this accurately, the regulator can infer that firm 1 is acting deceptively: under truthtelling, this distribution cannot arise as a conditional distribution of firm 1's reports, regardless of firm 2's type. In this manner, condition ($\bullet$) allows the regulator to identify any deceptions using conditional distributions of reports. In the next section, we describe a payment scheme that gives the firms' strict incentives to reveal these conditional distributions accurately.

3.2 The Mechanism

The mechanism is defined as follows. Each firm i reports a type $\hat{\theta}_i \in \Theta_i$ and a probability distribution $\hat{q}^i$ on Θ_{i-1}, which is to be interpreted as the distribution of the characteristics reported by $i - 1$. Pollution quantities are allocated as though the reported types of the firms are true: firm i is allocated $Q_i^*(\hat{\theta})$. Firm i's reported distribution is used in determining its own payment as well as firm $i-1$'s. In determining firm $i - 1$'s payment, the reported distribution $\hat{q}^i$ is compared to the distributions in the set $\{q_{\theta_i} \mid \theta_i \in \Theta_i\}$, the set of possible conditional distributions (depending on the realization of i's type) of firm $i - 1$'s reports *when $i - 1$ is reporting truthfully.* By reporting $\hat{q}^i$ not in this set, firm i indicates a deception by $i-1$. If $\hat{q}^i \in \{q_{\theta_i} \mid \theta_i \in \Theta_i\}$, then firm $i - 1$ pays $t^*(\hat{\theta})$; otherwise, firm $i - 1$ is paid

$$\left[\sum_{j \neq i-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j) \right] - C(Q^*(\hat{\theta})),$$

aligning firm $i - 1$'s incentives with those of the regulator in the manner of the mechanism of Groves (1973), Clarke (1971), and Vickrey (1960). We minimize their budget-balancing problem by employing their mechanism for out-of-equilibrium payments only.

To effectively use firm i's reported $\hat{q}^i$ to monitor firm $i-1$, firm i's reported distribution must also affect its own payment: we augment the above transfer to elicit truthful belief revelation as in Duggan (1998). Given $\hat{q}^i$ and reported characteristic $\hat{\theta}_{i-1}$, firm i pays

$$-\frac{\epsilon(1-\hat{q}^i(\hat{\theta}_{i-1}))^2}{2} + \sum_{\theta_{i-1}\neq\hat{\theta}_{i-1}} \frac{\epsilon\hat{q}^i(\theta_{i-1})^2}{2},$$

where $\epsilon > 0$ is arbitrary. We show below that this payment makes it a strict best response for every type θ_i to submit the actual conditional distribution of firm $i-1$'s reported types, regardless of $i-1$'s reporting strategy. Because we augment the firms' payments by negative amounts, the belief revelation part of the mechanism creates a net flow of transfers to the regulator, made arbitrarily small by choice of ϵ. We show below that, when there are at least three firms, payments can be modified to achieve full budget-balancing.

Proposition 2. *The unique pure strategy Bayesian equilibrium of the above mechanism is given by $(\hat{\theta}_i, \hat{q}^i) = (\theta_i, q_{\theta_i})$ for all i and all θ_i, and this yields the socially optimal allocation $Q^*(\theta)$ of pollution for all θ.*

We first argue that, in equilibrium, firm i correctly reports the distribution of firm $i-1$'s reported characteristics, regardless of $i-1$'s reporting strategy. Let $\hat{\theta}_{i-1}(\cdot)$ denote firm $i-1$'s (possibly deceptive) reporting rule, i.e., type θ_{i-1} of firm $i-1$ reports $\hat{\theta}_{i-1}(\theta_{i-1})$. Let r_{θ_i} denote the distribution of $i-1$'s reported characteristics conditional on θ_i. That is,

$$r_{\theta_i}(\theta'_{i-1}) \;=\; q_{\theta_i}(\{\theta_{i-1} \in \Theta_{i-1} \mid \hat{\theta}_{i-1}(\theta_{i-1}) = \theta'_{i-1}\}).$$

Given r_{θ_i} and reported distribution $\hat{q}^i$ for firm i, i's expected belief revelation payment to the regulator is

$$\sum_{\theta_{i-1}\in\Theta_{i-1}} \left[-\frac{\epsilon(1-\hat{q}^i(\theta_{i-1}))^2}{2} + \sum_{\theta'_{i-1}\neq\theta_{i-1}} \frac{\epsilon\hat{q}^i(\theta'_{i-1})^2}{2}\right] r_{\theta_i}(\theta_{i-1})$$

$$= \frac{\epsilon}{2}\sum_{\theta_{i-1}\in\Theta_{i-1}} \left[-(1-\hat{q}^i(\theta_{i-1}))^2 r_{\theta_i}(\theta_{i-1}) + \hat{q}^i(\theta_{i-1})^2(1 - r_{\theta_i}(\theta_{i-1}))\right].$$

This equality relies on

$$\sum_{\theta_{i-1}\in\Theta_{i-1}} \sum_{\theta'_{i-1}\neq\theta_{i-1}} \hat{q}^i(\theta'_{i-1})^2 r_{\theta_i}(\theta_{i-1}) \;=\; \sum_{\theta'_{i-1}\in\Theta_{i-1}} \hat{q}^i(\theta'_{i-1})^2(1 - r_{\theta_i}(\theta'_{i-1})),$$

which follows because, for all θ'_{i-1}, $\hat{q}^i(\theta'_{i-1})^2$ appears on the lefthand side once for each $\theta_{i-1} \neq \theta'_{i-1}$, where it is multiplied by $r_{\theta_i}(\theta_{i-1})$. For each θ_{i-1}, the unique solution to

$$\min_{0\leq q\leq 1} -\frac{(1-q)^2}{2}r_{\theta_i}(\theta_{i-1}) + \frac{q^2}{2}(1 - r_{\theta_i}(\theta_{i-1}))$$

13

is $r_{\theta_i}(\theta_{i-1})$. Thus, type θ_i of firm i has a strict best response to report the actual conditional probability distribution of $\hat{\theta}_{i-1}$ for *any* reporting rule used by firm $i-1$. In particular, if firm $i-1$ is reporting truthfully, type θ_i of firm i has a strict best response to report q_{θ_i}.

Now consider the strategy profile given in Proposition 2, which has each firm reporting its type truthfully and, conditional on its type, the true marginal distribution on its neighbor's type. We have shown that it is a best response for firm i to report the actual conditional distribution of firm $i-1$'s reported types, which in this case is just q_{θ_i}. Thus, the firms' reported distributions minimize belief revelation payments to the regulator. Since all firms, including firm $i+1$, are reporting true marginal distributions, firm i's payment is given by $t_i^*(\theta)$ for all θ and the additional payment to elicit beliefs. Then, since all other firms report their characteristics truthfully, the incentive compatibility of (Q^*, t^*) ensures that it is a best response for every type θ_i of firm i to report θ_i. Therefore, the strategy profile described in Proposition 2 is a Bayesian equilibrium.

To verify that it is the unique equilibrium, consider a pure strategy equilibrium where some firm i uses a deceptive reporting rule: $\hat{\theta}_i(\theta_i') \neq \theta_i'$ for some θ_i'. We have noted that every type $\tilde{\theta}_{i+1}$ of firm $i+1$ has a unique best response to report the actual conditional distribution $r_{\tilde{\theta}_{i+1}}$ of i's reported characteristics. We argue that $r_{\tilde{\theta}_{i+1}} \notin \{q_{\theta_{i+1}} \mid \theta_{i+1} \in \Theta_{i+1}\}$. Let $A = \{\theta_i'\}$ and $B = \{\theta_i \mid \hat{\theta}_i(\theta_i) = \theta_i'\}$, and note that, because type θ_i' is reporting falsely, $A \neq B$. Then, by condition ($\bullet$), $q_{\theta_{i+1}}(A) \neq q_{\tilde{\theta}_{i+1}}(B)$ for all θ_{i+1}. Equivalently, $q_{\theta_{i+1}}(\theta_i') \neq r_{\tilde{\theta}_{i+1}}(\theta_i')$ for all θ_{i+1}, as claimed.

Since the claim holds for every $\tilde{\theta}_{i+1}$, type θ_i of firm i's expected payment from the regulator is

$$\sum_{\theta_{-i} \in \Theta_{-i}} \left[\left[\sum_{j \neq i} B_j(Q^*(\hat{\theta}(\theta)), \hat{\theta}_j(\theta_j)) \right] - C(Q^*(\hat{\theta}(\theta))) \right] p_{\theta_i}(\theta_{-i}),$$

where $\hat{\theta}(\theta) = (\hat{\theta}_1(\theta_1), \ldots, \hat{\theta}_n(\theta_n))$ is the vector of reported characteristics, less the additional payment to elicit beliefs. Thus, firm i's expected payoff, conditional on characteristic θ_i, is

$$\sum_{\theta_{-i} \in \Theta_{-i}} \left[B_i(Q^*(\hat{\theta}(\theta)), \theta_i) + \left[\sum_{j \neq i} B_j(Q^*(\hat{\theta}(\theta)), \hat{\theta}_j(\theta_j)) \right] - C(Q^*(\hat{\theta}(\theta))) \right] p_{\theta_i}(\theta_{-i}),$$

less the expected belief revelation payment, which is independent of i's reported type. We analyze this expression term-by-term. Fix any θ_{-i}, let $\hat{\theta}_{-i} = \hat{\theta}_{-i}(\theta_{-i})$, and rewrite the θ_{-i} term of the expression as

$$\left[B_i(Q^*(\hat{\theta}_i, \hat{\theta}_{-i}), \theta_i) + \left[\sum_{j \neq i} B_j(Q^*(\hat{\theta}_i, \hat{\theta}_{-i}), \hat{\theta}_j) \right] - C(Q^*(\hat{\theta}_i, \hat{\theta}_{-i})) \right] p_{\theta_i}(\theta_{-i}).$$

14

By assumption, the unique maximum of $B_i(\cdot, \theta_i) + \left[\sum_{j \neq i} B_j(\cdot, \hat{\theta}_j)\right] - C(\cdot)$ is $Q^*(\theta_i, \hat{\theta}_{-i})$, which type θ_i can obtain by reporting θ_i. Moreover, since we have assumed that any false report would lead to a different pollution allocation, reporting θ_i is the unique maximizer. Since this is true for the term corresponding to every θ_{-i}, firm i's expected payoff, conditional on θ_i, is uniquely maximized by reporting the truth. But we began with the assumption that some type of firm i was reporting falsely in equilibrium, a contradiction.

Therefore, in any Bayesian equilibrium of the mechanism, every type of every firm must be reporting truthfully. The belief revelation payment ensures that every type θ_i of firm i reports the actual conditional distribution, q_{θ_i}, of firm $i - 1$'s reports. That is, $(\hat{\theta}_i, \hat{q}^i) = (\theta_i, q_{\theta_i})$ for all i and all θ_i.

3.3 Extensions

As before, our incomplete information results can be extended in a number of ways.

Non-generic Distributions of Characteristics. The arguments above depend critically on condition ($\bullet$). The condition implies that deceptions by firm i are necessarily reflected in the distribution of its reported characteristic, allowing the regulator to detect deceptions by firm i using firm $i + 1$'s reported $\hat{q}^{i+1}$. We have mentioned that the condition is generically satisfied but rules out some distributions of interest. We can relax condition ($\bullet$) by imposing the following *single-crossing property* on the benefit functions of the firms: let each firm i's characteristics be ordered $\theta_i^1 < \theta_i^2 < \cdots < \theta_i^{k_i}$, and assume that this ordering is preserved by the firm's marginal benefits, i.e.,

$$\frac{\partial B_i}{\partial Q_i}(Q_1, \ldots, Q_n, \theta_i^h) \quad < \quad \frac{\partial B_i}{\partial Q_i}(Q_1, \ldots, Q_n, \theta_i^{h+1})$$

for all $(Q_1, \ldots, Q_n)$. In words, firm i's higher types have higher marginal values of pollution. Furthermore, we require the optimal allocation of pollution to a firm to be monotonic in the firm's characteristic: $Q_i^*(\theta_i^h, \theta_{-i}) \leq Q_i^*(\theta_i^{h+1}, \theta_{-i})$ for all θ_{-i}. With these assumptions, we can drop the first part of condition ($\bullet$), leaving

($\circ$) For all i, all θ_i, and all $\theta_i' \neq \theta_i$, if A and B are distinct subsets of Θ_{i-1}, then $q_{\theta_i}(A) \neq q_{\theta_i'}(B)$.

Because the condition now applies only across distinct θ_i and θ_i', the distribution q_{θ_i} is restricted only in its relation to other conditional distributions. It could now be, for example, that q_{θ_i} is uniform for some θ_i.

Our uniqueness argument must now be modified. If q_{θ_i} is uniform, some deceptions by firm $i - 1$ (such as those where two types switch reports) will not be reflected in the distribution of $i - 1$'s reported types conditional on θ_i. We argue, however, that such deceptions cannot be best responses for firm $i - 1$. To illustrate, suppose that two types, $\theta_{i-1} < \theta_{i-1}'$, are switching reports. If it is a best response for type θ_{i-1} to

15

report θ'_{i-1}, thereby increasing its allotment of pollution, it cannot be a best response for type θ'_{i-1}, which values this increment of pollution more, to report θ_{i-1}. Generally, if it is desirable for one type θ_{i-1} to report falsely in one direction (say, higher), then it cannot be a best response for higher types to report θ_{i-1} or lower. Therefore, any deception that could be a best response for $i-1$ will be reflected in the firm's distribution of reported characteristics, and our above analysis applies.

Mixed Strategies. Allowing firm $i-1$ to use mixed reporting strategies raises the possibility of a mixed deception that is not reflected in the conditional distributions of $i-1$'s reported characteristics. In 2×2 example in the previous subsection, suppose that type θ''_1 reports θ'_1 with probability one, and that type θ'_1 reports the truth with probability $1/2$ and θ''_1 with probability $1/2$. Consider the distribution of firm 1's reports conditional on type θ'_2 of firm 2: firm 2 expects firm 1 to report θ'_1 with probability $2/3$ and θ''_1 with probability $1/3$. In contrast to the pure strategy case, the conditional distribution induced by this mixed deception is consistent with truthtelling.

Such examples of consistent mixed deceptions require an exact balancing of reporting probabilities: if some type θ_i reports deceptively with positive probability, other types must report θ_i to just offset that probability. Ordering the characteristics of the firms and imposing the single-crossing property defined above, our previous arguments show that this balancing cannot be a best response for all of a firm's types, and therefore cannot occur in equilibrium. Thus, under the single-crossing assumption, our arguments can be extended to allow firms to use mixed strategies.

Budget-balancing. When $n \geq 3$, the mechanism is easily modified to achieve budget-balancing in equilibrium, assuming t^* is balanced. We simply make firm $i+2$'s belief revelation payment a transfer to firm i. That is, we deduct the amount

$$-\frac{\epsilon(1-\hat{q}^{i+2}(\hat{\theta}_{i+1}))^2}{2} + \sum_{\theta_{i+1} \neq \hat{\theta}_{i+1}} \frac{\epsilon \hat{q}^{i+2}(\theta_{i+1})^2}{2},$$

depending on the reports of $i+1$ and $i+2$, from firm i's payment to the regulator. Because the added payment does not depend on the reports of firm i, the above analysis is not affected.

By further modifying the mechanism, we can achieve budget-balance out of equilibrium as well. We continue to have each firm i report $\hat{\theta}_i$ and $\hat{q}^i$, with belief revelation payments specified as above. If exactly one firm, say i, reports $\hat{q}^i$ inconsistent with truth-telling by its neighbor, firm $i-1$ is paid

$$\left[\sum_{j \neq i-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j)\right] - C(Q^*(\hat{\theta}));$$

16

each firm $j \neq i - 1$ pays $t_j^*(\hat{\theta})$; and, letting l be the highest indexed firm other than i or $i - 1$, in addition firm l pays

$$\left[\sum_{j \neq i-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j) \right] - C(Q^*(\hat{\theta})) - t_{i-1}^*(\hat{\theta}),$$

balancing the budget. If more than one firm reports a distribution inconsistent with truth-telling by its neighbor, let i and k be the lowest indexed such firms; then firm $i - 1$ is paid

$$\left[\sum_{j \neq i-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j) \right] - C(Q^*(\hat{\theta}));$$

firm $k - 1$'s payment is defined similarly, with "$k - 1$" replacing "$i - 1$"; each firm j other than $i - 1$ and $k - 1$ pays $t_j^*(\hat{\theta})$; and, letting l denote the highest indexed firm other than $i - 1$ and $k - 1$, in addition firm l pays

$$\left[\sum_{j \neq i-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j) \right] + \left[\sum_{j \neq k-1} B_j(Q^*(\hat{\theta}), \hat{\theta}_j) \right] - 2C(Q^*(\hat{\theta})) - t_{i-1}^*(\hat{\theta}) - t_{k-1}^*(\hat{\theta}),$$

balancing the budget. The strategy profile of Proposition 2 is still clearly an equilibrium: if, for example, some type θ_i of firm i deviates by reporting $\hat{q}^i \neq q_{\theta_i}$, its own payments to the regulator are unchanged.

To see that there are no other equilibria, consider any equilibrium strategy profile. We first claim that firm 1 reports the actual distribution of firm n's reported characteristics. We have shown that firm 1's expected belief revelation payment to the regulator, conditional on each type, is minimized by doing so, but we must verify that doing so will not affect the remainder of firm 1's payment. Consider any realization of the other firms' types and any reports by the other firms. If no other firms report inconsistencies, then firm 1 pays $t_1^*(\hat{\theta})$ regardless of its reported distribution. If some other firm, say i, reports a distribution inconsistent with truth-telling by its neighbor, then either $i - 1 = 1$, or firm 1 is used to balance the budget, or neither holds. In the first case, firm 1 receives the Groves-Clarke-Vickrey transfer regardless of its reported distribution. The second case can only occur when $n = 3$, firm 2 does not report an inconsistency with truth-telling on the part of firm 1, and firm 3 does report an inconsistency on the part of firm 2. In that case, firm 1 balances the budget regardless of its reported distribution. If neither holds, then firm 1 receives $t_1^*(\hat{\theta})$ regardless of its reported distribution. Thus, the claim is proved.

We next claim that firm n reports its characteristic truthfully. If not, we have just established that each type of firm 1 will report a distribution for n inconsistent with truth-telling. Then, by construction, firm n receives the Groves-Clarke-Vickrey transfer with probability one and, as explained in Section 3.2, firm n has a strict incentive to report its characteristic truthfully, a contradiction. Furthermore, it follows that each type of firm 1 reports a distribution for n consistent with truth-telling.

A similar argument establishes that firm 2 accurately reports the distribution of firm 1's reported characteristic. The only point that may require elaboration is when some other firm, i, reports an inconsistency and firm 2 is used to balance the budget. Since firm 1 does not report an inconsistency, this case can only occur when $n = 4$, firm 3 does not report an inconsistency on the part of firm 2, and firm 4 does report an inconsistency on the part of firm 3. In that case, firm 2 balances the budget regardless of its reported distribution. Thus, firm 2 accurately reports the distribution of firm 1's reported characteristic, and it follows as above that firm 1 must report its own characteristic truthfully and that each type of firm 2 reports a distribution for 1 consistent with truth-telling. Arranging these arguments in the form of an induction proof (omitted), we see that the same is true for all firms and, consequently, that the equilibrium is exactly that of Proposition 2.

References

[1] E. Clarke (1971) "Multipart Pricing of Public Goods," *Public Choice,* 8: 19-33.

[2] J. Crémer and R. McLean (1985) "Optimal Selling Strategies under Uncertainty for a Discriminating Monopolist when Demands Are Interdependent," *Econometrica,* 53: 345-361.

[3] J. Crémer and R. McLean (1988) "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica,* 56: 1247-1257.

[4] C. d'Aspremont and L.A. Gerard-Varet (1979)"On Bayesian Incentive Compatible Mechanisms," in Jean-Jacques Laffont, ed., *Aggregation and Revelation of Preferences,* North-Holland Publishers.

[5] J. Dales (1968) *Pollution, Property, and Prices,* Toronto: University of Toronto Press.

[6] P. Dasgupta, P. Hammond, and E. Maskin (1980) "On Imperfect Information and Optimal Pollution Control," *Review of Economic Studies,* 67: 857-860.

[7] J. Duggan (1998) "An Extensive Form Solution to the Adverse Selection Problem in Principal/multi-agent Environments," *Review of Economic Design,* 3: 167-191.

[8] J. Duggan and J. Roberts (1998) "Robust Implementation," mimeo.

[9] J. Eyckmans (1997) "Nash Implementation of a Proportional Solution to International Pollution Control Problems," *Journal of Environmental Economics and Management,* 33: 314-30.

[10] T. Groves (1973) "Incentives in Teams," *Econometrica,* 41: 617-631.

[11] T. Groves (1979) "Efficient Collective Choice when Compensation Is Possible," *Review of Economic Studies,* 46: 227-241.

[12] R. Hahn (1984) "Market Power and Transferable Property Rights," *Quarterly Journal of Economics,* 99: 753-765.

[13] R. Hahn and G. Hester (1989) "Where Did All the Markets Go? An Analysis of EPA's Emissions Trading Program," *Yale Journal on Regulation,* 6: 109-153.

[14] J. Harsanyi (1967-68) "Games with Incomplete Information Played by 'Bayesian' Players," *Management Science,* 14: 159-182, 320-334, 486-502.

[15] J. Herriges, R. Govindasamy, and J. Shogren (1994) "Budget-Balancing Incentive Mechanisms," *Journal of Environmental Economics and Management,* 27: 275-85.

[16] L. Hurwicz (1979) "Outcome Functions Yielding Walrasian and Lindahl Allocations at Nash Equilibrium Points," *Review of Economic Studies,* 46: 217-225.

[17] M. Jackson (1991) "Bayesian Implementation," *Econometrica,* 59: 461-477.

[18] A. Kritikos (1993) "Environmental Policy under Imperfect Information: Comment," *Journal of Environmental Economics and Management,* 25: 89-92.

[19] E. Kwerel (1977) "To Tell the Truth: Imperfect Information and Optimal Pollution Control," *Review of Economic Studies,* 44: 595-601.

[20] T. Lewis and D. Sappington (1995) "Using Markets to Allocate Pollution Permits and Other Scarce Resource Rights Under Limited Information," *Journal of Public Economics,* 57: 431-455.

[21] P. McAfee and P. Reny (1992) "Correlated Information and Mechanism Design," *Econometrica,* 60: 395-421.

[22] E. Maskin (1977) "Nash Implementation and Welfare Optimality," mimeo.

[23] E. Maskin (1979) "Implementation and Strong Nash Equilibrium," in Jean-Jacques Laffont, ed., *Aggregation and Revelation of Preferences,* North-Holland Publishers.

[24] H. Matsushima (1993) "Bayesian Monotonicity with Side Payments," *Journal of Economic Theory,* 59: 107-121.

[25] J. Moore and R. Repullo (1988) "Subgame Perfect Implementation," *Econometrica,* 56: 1191-1220.

[26] J. Moore and R. Repullo (1990) "Nash Implementation: A Full Characterization," *Econometrica,* 58: 1083-1099.

[27] R. Sundaram (1996) *A First Course in Optimization Theory,* Cambridge: Cambridge University Press.

[28] H. Varian (1994) "A Solution to the Problem of Externalities When Agents are Well-informed," *American Economic Review,* 85: 1278-1293.

[29] W. Vickrey (1960) "Counterspeculation, Auctions, and Competitive Sealed Tenders," *Journal of Finance,* 16: 8-37.

[30] M. Walker (1981) "A Simple Incentive Compatible Scheme for Attaining Lindahl Allocations," *Econometrica,* 49: 65-71.

[31] A. Xepapadeas (1991) "Environmental Policy under Imperfect Information," *Journal of Environmental Economics and Management*, 20: 113-26.