A Simple Auction Mechanism for the Optimal Allocation of the Commons

Juan-Pablo Montero*

November 20, 2006

Abstract

Efficient regulation of the commons requires information about the regulated firms that is rarely available to regulators (e.g., cost of pollution abatement). Different mechanisms have been proposed for inducing firms to reveal their private information but, for reasons discussed in the paper, these mechanisms are of limited use. This paper proposes a much simpler mechanism that implements the first-best for any number of firms: a uniform price sealed-bid auction of an endogenous number of (transferable) licenses with a fraction of the auction revenues given back to firms. Paybacks, which rapidly decrease with the number of firms, are such that truth-telling is a dominant strategy regardless of whether firms behave non-cooperatively or collusively. (*JEL* D44, D62, D82)

Regulatory authorities generally find that part of the information they need for implementing an efficient regulation is in the hands of those who are to be regulated. Regulating externalities such as access to common resources (e.g., clean air, water streams, fisheries, native forests, etc.) is not the exception. Environmental regulators, for example, know little about firms' pollution abatement costs, so without communicating with firms they would be unable to establish the efficient level of pollution. Different mechanisms have been proposed for inducing firms to reveal their private information but for reasons I discuss below I find these mechanisms

*Department of Economics, Pontificia Universidad Católica de Chile, Vicuña Mackenna 4860, Santiago, Chile, and MIT's Center for Energy and Environmental Policy Research (jmontero@faceapuc.cl). I would like to thank Tony Creane, Evan Kwerel, Matti Liski, Marty Weitzman, three anonymous referees, seminar participants at Harvard University, Texas A&M University, University of Maryland, University of Massachusetts-Amherst, Universidade de Vigo and Universidad Alberto Hurtado, and specially Bill Hogan for many reactions and comments. Most part of the article was written while I was visiting Harvard's Kennedy School of Government (KSG) under a Repsol YPF-KSG Research Fellowship. Financial support from Fondecyt-Chile is also gratefully acknowledged. All errors and omissions are mine.

of limited use. In this paper, I propose an efficient and much simpler mechanism: a uniform price sealed-bid auction of an endogenous number of (transferable) licenses with a fraction of the auction revenues given back to firms.[1] The mechanism is developed under the additional assumption that firms know nothing about the other firms' characteristics.

Following Martin L. Weitzman (1974), several authors have looked for ways in which to improve upon his fixed tax or license scheme. Marc J. Roberts and Michael Spence's (1976) hybrid tax/license scheme can in principle implement the first-best when there is an infinitely large number of firms and the regulator is free to impose a tax schedule (as opposed to a fixed tax) and issue a continuum of license types, with each type clearing at a different price. Building also upon the assumption of perfect competition in the license market, Evan Kwerel (1977) develops a simpler subsidy/license scheme that implements the first-best in Nash equilibrium. Relaxing the perfect competition assumption and allowing for pollutant differentiation, Partha Dasgupta, Peter Hammond and Eric Maskin (1980) propose a simple adaptation of the Vickrey-Clarke-Groves (VCG) mechanism, which has the advantage of implementing the first-best in dominant strategies.[2] Maintaining the same two assumptions but imposing complete information by firms, in a more recent paper John Duggan and Joanne Roberts (2002) advance an efficient quantity-based scheme in which each firm chooses the number of licenses for itself and for its "neighbor."[3]

With the exception of Kwerel (1977), the fact that these first-best mechanisms are highly non-linear, firm-specific and/or depending on a complete information assumption complicates their practical implementation.[4] In fact, we do not see anything like it being applied in practice or, at least, under consideration. Conversely, Kwerel's scheme is quite simple: the regulator issues a fixed number of transferable licenses and establishes a subsidy per license to be paid to any firm holding licenses in excess of its emissions. Both the total number of licenses and the subsidy level are calculated on the basis of the information provided by firms. Unfortunately, the scheme presents some important limitations regardless of how licenses are allocated to firms.[5] If licenses are allocated for free (i.e., grandfathered), it can be shown (Proposition 1)

[1] Licenses are generally refered to as permits or allowances in water and air pollution control, as rights in water supply management and as quotas in fisheries management. In this paper, I will use the term license throughout.

[2] See Daniel F. Spulber (1988) for consideration of budget constraints.

[3] Other efficient proposals that require complete information by firms (whether globally or at the industry level) are Hal R. Varian's (1994) multistage price mechanism, Jae-Cheol Kim and Ki-Bok Chang's (1993) tax/subsidy scheme and Marcelo Caffera and Juan Dubra's (2006) industry-specific emission standard mechanism.

[4] Evidence of significant information asymmetries across firms facing a commons problem is provided, for example, by Steven N. Wiggins and Gary D. Libecap (1985).

[5] Kwerel (1977) is never explicit on the allocation of licenses other than assuming that firms pay a uniform

that firms have incentives not to reveal their true types but to over-report their demand for licenses to the maximum extent possible.

If, on the other hand, the total number of licenses are allocated via a uniform-price auction in which each firm bids a demand schedule indicating the number of licenses willing to purchase at any given price, there is little guarantee that firms will reach the competitive outcome (Paul Milgrom, 2004). As first recognized by Robert Wilson (1979) in his pioneer "auctions of shares" article, even when there is a large number of bidders, uniform price auctions can exhibit Nash equilibria with prices far below the competitive price (the price that would prevail if all bidders submit their true demand curves). The reason for this is that uniform pricing creates strong incentives for bidders to (non-cooperatively) under-report their demand schedules in order to depress the price they pay for their inframarginal units. Therefore, anticipating a low-price equilibrium at the auction (most likely zero, or the reserve price if there is any), firms would find it again profitable to over-report their demand functions by as much as possible in order to acquire a large volume of licences that can then be sold back to the government at a price much higher than the auction clearing price.[6]

Different solutions have been advanced in dealing with this low-price equilibria phenomenon. One radical solution is to give up the uniform-price format altogether and opt for a discriminatory-price format (e.g., William Vickrey, 1961; Lawrence M. Ausubel, 2004). But within the uniform-price format, different authors have also been looking for ways in which changing auction rules could eliminate underpricing. Ilan Kremer and Kjell G. Nyborg (2004), for example, propose changing the allocation rule (i.e., the way the asset is divided when there is excess demand at the clearing price) from the usual marginal pro rata tie-breaking rule to a total pro rata rule.[7] More recently, David McAdams (2005) eliminates underpricing by letting the auctioneer not to commit to a fixed quantity and reserve price ex-ante. Bidders only learn

price for all licenses purchased. Besides, there is no much we can infer from the firm's cost minimizing problems laid out in pages 596 and 597 because of the price-taking behavior (i.e., any grandfathered allocation can be omitted from the minimization problem since it is a lump-sum transfer with no effect on the firm's abatement decision). It is also not obvious how to adapt Kwerel's scheme to the case of pollution differentiation.

[6] Some readers may argue that removing the subsidy and having it replaced by a price floor (of equal magnitude) would finally solve matters. As we shall see, it does not do so for the exact opposite reasons. Anticipating the reserve price as the auction clearing price, firms would have incentives to deviate from truth-telling by under-reporting their demand curves in order to decrease the reserve price. The expectation of a (non-cooperative) low-price equilibrium acts as coordinating monopsonist device.

[7] For the change in the allocation rule to have an effect on clearing prices, bidders must be allowed to submit discountinuous demand schedules, which, by construction, is not possible in Wilson (1979). But unlike in Kremer and Nyborg (2004) where bidders have a constant valuation for the asset, in our context this allocation rule change is of little help because bidders do have fairly continuos downward sloping demand curves.

about the total quantity sold by the auctioneer once the auction is concluded.

In this paper I propose a mechanism that builds upon a conventional uniform price sealed-bid auction but in order to guarantee the efficient outcome, I introduce two key ingredients. First, I let the total number of licenses be endogenous to the demand schedules submitted by firms. This is a most natural thing to do in our context because the benevolent regulator is clueless about the efficient number of licenses to be allocated before communicating with firms. But unlike in McAdams (2005), this "flexible supply" feature by itself does not fully solve the underpricing problem.[8] Hence, I introduce a second ingredient: rebates or paybacks. Part of the auction revenues are returned to firms not as lump sum transfers but in a way that firms would have incentives to bid truthfully. While rebates may seem odd in other contexts,[9] they are not new in existing auctions for "protecting the commons".[10] Furthermore, an auction with paybacks seem to be a natural point of departure for any license-type regulatory proposal given the mixed experience with allocating licenses (grandfather allocation vs auction allocation) that is observed in existing programs across a variety of areas including air-pollution control, water supply management and fisheries management (Tom Tietenberg, 2003).[11]

The two ingredients —endogenous supply of licenses and paybacks— enter into the uniform-price format in a way that the resulting auction mechanism is both ex-post efficient and strategy-proof (i.e., implemented in dominant strategies).[12] The supply curve of licenses reflects the cost to society (other than firms) from allocating these licenses to firms. Paybacks, on the other hand, are such that the total payment for licenses of each firm is exactly equal to the "damage" it exerts upon all the other agents (i.e., other regulated firms and the rest of society). For example, in the case of a single polluting firm, the total payment faced by the firm is equal to the pollution damage $D(l)$, where l is the number of licenses/pollution allocated to the firm

[8] It only works in the limit, when there is an infinitely large number of firms so that paybacks are virtually zero. Part of the reason why "flexible supply" is not sufficient is because I work with a very different set of assumptions than McAdams (2005). I let firms to be asymmetric, to have downward sloping demand curves and to know nothing about other firms' characteristics. In addition, my auctioneer's objective function is not to maximize revenues but social welfare.

[9] Not surprisingly they are absent in recent books by Milgrom (2004) and Paul Klemperer (2004).

[10] See, for example, the US EPA auction for sulfur dioxide allowances (Paul L. Joskow et al., 1998). See also Hans Gersbach and Till Requate (2004).

[11] See also Peter Cramton and Suzi Kerr (2002) for related arguments.

[12] This result may look surprising at first, given the result of Jerry R. Green and Jean-Jacques Laffont (1979) that there is no social choice function that is strategy-proof and ex-post efficient. That result relies crucially on the fact that the type of all the players is not known. In this paper, however, the type of one player, the benevolent regulator, is known, that is, the "mechanism designer" knows the regulator's loss function from allocating licenses to firms (e.g., the pollution damage function). Given that, it is well known that using the known player as a money sink or source one can design an ex-post efficient and strategy-proof mechanism (Andreu Mas-Colell et al., 1995, pp. 876-82).

at the auction. In the case of multiple firms, the total payment faced by firm i is equal to its residual damage $D_i(l_i)$, which involves both its pecuniary externality imposed upon other competing bidders (i.e., regulated firms) and its residual pollution externality. The residual damage function $D_i(l_i)$ is computed by substracting from the supply curve $D'(l)$ all other firms' bids, so it is independent of firm i's bid.

The auction mechanism certainly follows a VCG payoff rule in that it makes each firm to pay exactly for the externality it imposes on the other agents.[13] It is, however, structurally different from the VCG mechanism of Dasgupta-Hammond-Maskin (DHM). For instance, the auction mechanism does not explicitly exhibit the VCG constant term constructed upon other firms' reports. Because of the structural differences the two mechanisms differ in at least two important ways. First, the DHM mechanism, unlike the auction mechanism, fails to allocate licenses efficiently across firms when the aggregate supply of licenses is fixed. This is because each individual firm is no longer pivotal under DHM, i.e., its report does not affect the aggregate supply. This is quite an important distinction because in many commons problems the aggregate supply is likely to be fixed, either because of the presence of some genuine threshold or because the auctioneer/regulator has no control over the aggregate supply. Second, the DHM mechanism, unlike the auction mechanism, fails to deliver the first-best when firms are acting collusively unless the constant term is set to zero. When we do that, collusive and non-cooperative behavior are indeed no different in DHM but payments become so large —the full social cost— that the (now Groves) mechanism becomes of little practical value.

The article is organized as follows. Section I presents the modelling assumptions and a brief discussion of Kwerel's scheme. Section II describes the auction mechanism; first for a single firm and then for multiple firms. Section III describes several properties of the mechanism including its relationship to the VCG mechanism. Section IV looks at how the mechanism performs under collusive behavior. Section V concludes. The Appendix contains the proofs to all propositions.

I. The Model

To facilitate the exposition I will develop the model for the case of a classical pollution externality (which would correspond to an auction of shares with variable supply). But it is worth emphasizing that the model readily extends to other commons problems including those in which licenses are firm-specific (because of pollution differentiation) and where firms impose

[13] It is also payoff equivalent to the discriminatory auctions of Vickrey (1961) and Ausubel (2004).

(private) externalities on each other.[14]

A. *Notation and First-Best Allocation*

Consider $n \geq 1$ firms ($i = 1, ..., n$) to be regulated. All firms are assumed to have inverse demand functions for pollution of the form $P_i(x_i)$ with $P_i'(x_i) < 0$, where x_i is firm i's pollution level that is accurately monitored by the regulator (In some cases I will work with the demand function, which is denoted by $X_i(p)$ with $X_i'(p) < 0$, where p is the price of pollution). Function $P_i(\cdot)$ is only known by firm i, neither by the regulator nor by the other firms. The aggregate demand curve for pollution is denoted by $P(x)$, where $x = \sum_{i=1}^{n} x_i$ is total pollution. The social damage caused by pollution x is $D(x)$ with $D(0) = 0$, $D'(x) > 0$ and $D''(x) \geq 0$. $D'(x)$ can be interpreted more generally as the regulator's supply function for licenses. We may want to assume that $D(x)$ is publicly known but it is actually not necessary.

In the absence of regulation firm i would emit x_i^0 , where $P_i(x_i^0) = 0$. Hence, firm i's cost of reducing emissions from x_i^0 to some level $x_i < x_i^0$ is $C_i(x_i) = \int_{x_i}^{x_i^0} P_i(z)dz$ (note that $-C_i'(x_i) \equiv P_i(x_i)$), and the minimum total cost of achieving pollution level $x < x^0$ is $C(x) = \int_{x}^{x^0} P(z)dz$.

The regulator's objective is to minimize the sum of clean-up costs and damages from pollution, i.e., $C(x) + D(x)$. Therefore, the socially optimal or first-best pollution level $x^* < x^0$ satisfies

$$P(x^*) = D'(x^*) = P_i(x_i^*) \quad \text{for all } i = 1, ..., n \tag{1}$$

But the regulator cannot directly implement the first-best allocation because he does not know the demand functions $P_i(\cdot)$. He must then look for mechanisms in which it is in the firms' best interest to communicate their private information to him. Kwerel (1977) advances one of such mechanisms for the case in which there are many firms.

B. *Kwerel's Scheme*

To appreciate the workings of my auction scheme it is useful to start by understanding firms' incentives under Kwerel's scheme. This latter proves to be interesting in itself because, as we shall see below, the scheme may not work as intended.

Kwerel's mechanism is based on the combination of two instruments: an allocation of a total of l transferable licenses and a subsidy of s per license to be paid to any firm holding

[14] Formal analysis of these extensions are found in Montero (2006).

licenses in excess of its emissions. The regulator asks firms to report their demand curves after they are informed that the parameters l are s are to be set according to

$$s = \hat{P}(l) = D'(l) \tag{2}$$

where $\hat{P}(\cdot)$ is the aggregate demand curve built upon individual reports $\hat{P}_i(x_i)$.

Assuming that the market for licenses is competitive, it must hold in equilibrium that $P_i(x_i) \equiv -C_i'(x_i) = p$ and $x_i = l_i$ for all $i = 1, ..., n$, where p denote the market price of licenses. Firms equate marginal abatement costs to the market price and keep a number of licenses just to cover their emissions. Kwerel argues that this simple scheme induces each firm i to report its true demand curve $P_i(\cdot)$ as long as it believes all other firms are telling the truth. In other words, truth-telling is a (Bayesian) Nash equilibrium.

Kwerel's logic can be easily explained with the aid of Figure 1. Figure 1a depicts the situation in which a firm or a group of firms over-report their demand curves such that the reported aggregate demand curve is $\hat{P}(x)$ instead of the true curve $P(x)$. The license and subsidy parameters take the values of $\hat{l}$ and $\hat{s}$, respectively, which are above their first-best levels l^* and s^*. Since the government is buying back licenses at price $\hat{s}$, the market equilibrium price of licenses is not p' (as if no license were sold back to the government) but $p = \hat{s} > p^*$. On aggregate, firms sell back $\hat{l} - \hat{x}$ licenses, so total pollution falls below its first-best level to $\hat{x} < l^*$. Figure 1b, on the other hand, depicts the under-reporting situation. Given the reported aggregate demand curve $\tilde{P}(x)$, the license and subsidy parameters take now the values of $\tilde{l}$ and $\tilde{s}$, respectively. The market equilibrium price is $p = \tilde{p} > p^*$ and total pollution is $x = \tilde{l} < l^*$.

From inspection of these two cases it should become evident that no matter what firms report to the regulator the market equilibrium price of licenses is given by $p = \max\{P(x), D'(x)\}$. Hence, the minimum possible equilibrium price for licenses is $p^{\min} = P(x^*) = D'(x^*)$, which is obtained when all firms report their true demand curves. Based on this observation, Kwerel closes his proof by arguing that since each firm's compliance cost is an increasing function of p, no firm has incentives to move the aggregate demand curve from its actual value, whatever it is, when it believes that all the other firms are telling the truth.[15]

Kwerel's logic holds as long as all licenses are allocated via a competitive (uniform-price)

[15] Kwerel also mentions the existence of multiple "offsetting-lies" Nash equilibria in which two or more firms send false reports that, on aggregate, add to the true demand curve $P(x)$. Without knowing the actual $P(x)$, however, it is hard to see how firms could coordinate in one of these "offsetting-lies" Nash equilibria.

auction. But we already discussed that auctions of shares are rarely competitive because of under-reporting incentives. If, on the other hand, licenses are allocated for free, the revenues accruing to firms from selling licenses back to the government create over-reporting incentives. To see this, simply go back to Figure 1a and compare the total compliance costs from reporting the true aggregate demand curve $P(x)$, i.e., area $x^0\hat{x}A$, with those from reporting $\hat{P}(x)$, i.e., area $x^0 l^* E$ minus area $\hat{l}\hat{x}AB$. Clearly, the regulation has turned out to be quite a profitable business for firms, and more so the higher the degree of over-reporting. More generally, it can be established

PROPOSITION 1: *The unique (Nash-equilibrium) outcome in Kwerel's scheme under a free allocation of licenses is for firms to over-report their demand curves as to ensure the maximum possible number of licenses and subsidy level.*

Despite its limitations, Kwerel's scheme has an element that I also use in constructing the auction mechanism that I present next. Under this new scheme firms are also communicated in advance that the information they report to the regulator will be used in a form similar to expression (2); although with some fundamental differences.

II. The Auction Mechanism

It helps to start with the single-firm case. I will then extend the mechanism to the general case of multiple (non-cooperative) firms.

A. *Single Firm*

Consider a single firm with demand curve $P(x) \equiv -C'(x)$. The auction scheme operates as follows. First, the firm is informed in advance about the auction rules (including the way the auction clears and the paybacks are computed). Then, the firm is asked to bid a non-increasing inverse demand schedule $\hat{P}(x)$ (or equivalently, a non-increasing demand schedule $\hat{X}(p)$). With this information, the auctioneer/regulator clears the auction (i.e., determines p and l) according to

$$p = \hat{P}(l) = D'(l) \tag{3}$$

The firm receives l licenses and pay p for each license. Soon after the firm gets a fraction $\alpha(l)$ of the auction revenues back (i.e., payback is $\alpha(l)pl$).

It is readily seen in Figures 1a and 1b that it is not socially optimal for the regulator to set the fraction $\alpha(l)$ equal to either 1 or 0. If the regulator keeps no revenue for himself (i.e.,

$\alpha(l) = 1$), the firm has incentives to over-report by as much as to postpone any abatement effort. Conversely, if the regulator keeps all the auction revenues for himself (i.e., $\alpha(l) = 0$), the firm has incentives to under-report to some optimal extent. By submitting $\tilde{P}(x)$ instead of $P(x)$ in Figure 1b, the firm is able to reduce its compliance cost from area $x^0 0 p^* E$ to area $x^0 0 \tilde{s} F B E$. The firm's optimal under-reporting in this case balances at the margin the gains from getting a lower price for licenses with the losses from higher abatement.

To find the function $\alpha(l)$ that just induces the firm to submit its true demand curve and hence allows the regulator to implement the first-best, we proceed by backward induction. Given some function $\alpha(l)$, the firm's problem is to find the demand schedule $\hat{P}(x)$ that solves

$$\min C(l) + pl - \alpha(l)pl \tag{4}$$

subject to (3).

Using the auction clearing equation (3) we can replace p by $D'(l)$ in (4), and since there is a one-to-one correspondence between a demand schedule $\hat{P}(x)$ and the number of licenses l, at least in the range of prices the firm expects the auction to clear, the firm's first order condition is given by

$$C'(l) + D'(l) + D''(l)l - \alpha'(l)D'(l)l - \alpha(l)(D''(l)l + D'(l)) = 0 \tag{5}$$

Anticipating (5), the regulator's problem is to find the function $\alpha(l)$ that induces the firm to deliver the first-best allocation, i.e., $C'(l^*) + D'(l^*) = 0$ (or $P(l^*) = D'(l^*)$). Such function solves the differential equation

$$\alpha'(l) + \alpha(l)\left(\frac{D''(l)l + D'(l)}{D'(l)l}\right) = \frac{D''(l)}{D'(l)} \tag{6}$$

The function $\alpha(l)$ that results from solving (6) is the function the regulator informs the firm along with the other auction rules. But if $\alpha(l)$ is such that the firm is delivering the first-best l^* it must be the case that the firm is solving the regulator's problem up front. In other words, the last two terms of (4) must add to $D(l)$, which leads to

PROPOSITION 2: *The payback function is given by*[16]

$$\alpha(l) = 1 - \frac{D(l)}{D'(l)l}$$

[16] *Strictly speaking the solution of (6) is* $\alpha(l) = 1 - D(l)/D'(l)l + K/D'(l)l$, *where* K *is an integration constant that represents –unless it is set to zero– a lump-sum transfer from(to) the firm.*

Since $D'(l)$ is a non-decreasing function of l it is clear that $0 \leq \alpha(l) \leq 1$, so the final price paid by the firm for each license (i.e., $(1-\alpha)p$), is at most equal to marginal damage $D'(l)$. Plugging back the function $\alpha(l)$ of Proposition 2 into the firm's objective function (4) it is immediately seen that the new auction scheme has indeed converted the firm's problem into the regulator's by making the firm bear the full cost of the pollution damages.

The idea of requiring the firm to pay $D(l)$ is certainly not new. The DHM mechanism for the case of a single firm reduces precisely to the regulator informing the firm that it faces a tax function $T(x) = D(x)$, where x is the firm's observed pollution level (note also that because there is only one firm the regulator does not need the firm to report any cost/demand information to him). The auction mechanism implements the same result but in a slightly different way. Here the regulator asks the firm to submit a demand schedule that it is then used to compute the optimal number of licenses and the price to be charged for each license. In that sense the auction scheme fully decouples the regulatory design from the enforcement/monitoring activity like in any other quantity-based regulation —whether it is based on standards or transferable licenses— which today is by far the more prevalent type of regulation for protecting the commons.[17] Other differences will become evident as we consider more than one firm.

But before moving to the multiple-firm case it is worth mentioning that if our single firm knows the function $D(x)$, it does not need to truthfully bid its entire demand schedule but only the portion relevant to the auction clearing. It could for instance submit the perfectly inelastic demand schedule $\hat{X}(p) = l^*$. More importantly, although I have developed the auction mechanism as if the firm knew the damage function $D(x)$, it should be clear by now that the firm does not actually need to know $D(x)$, and hence $\alpha(l)$, for the auction mechanism to work in providing incentives for truthful revelation. We only require the firm to believe that it is facing a regulator committed to implement the first-best for whatever function $D(x)$ he has in mind. And if the firm does indeed know little about $D(x)$, it will truthfully bid its (almost) entire demand schedule to make sure that for any possible function $D(x)$ chosen by the regulator it will get the first-best level of licenses.

B. *Multiple Firms*

Let us now consider the incomplete information game of $n \geq 2$ firms with $P_i(x_i)$ being firm

[17] It is easy, however, to transform the DHM mechanism into a quantity-based mechanism by simply allocating licenses at the reported first-best levels and asking firms to pay for them according to their individual tax schedules.

i's private information. The auction mechanism extends as follows. Firm i $(= 1, ..., n)$ is asked to bid a non-increasing inverse demand schedule $\hat{P}_i(x_i)$ (or, equivalently, a non-increasing demand schedule $\hat{X}_i(p)$). Based on this information, the regulator computes the residual marginal damage function (i.e., residual supply function) for each firm i using other firms' demand schedules, which is (see Figure 2)

$$D_i'(x_i) \equiv D'(x) - \hat{P}_{-i}(x_{-i}) \tag{7}$$

where $\hat{P}_{-i}(x_{-i})$ is the demand schedule from aggregating all other bids but firm i's and $x_i \equiv x - x_{-i} \geq 0$ (note that since $x_i \geq 0$, $D_i'(x_i)$ is only defined at and above the point at which $D'(x) = \hat{P}_{-i}(x_{-i})$; see Figure 2). The regulator clears the auction by determining a price p_i and number licenses l_i for each bidder i according to

$$p_i = \hat{P}_i(l_i) = D_i'(l_i)$$

Firm i purchases l_i licenses at a price p_i each, and soon after gets a payback of $\alpha_i(l_i)p_i l_i$, where $\alpha_i(l_i)$ is the payback fraction specific to firm i.

If $\alpha_i(l_i)$ is set according to Proposition 2, i.e.,

$$\alpha_i(l_i) = 1 - \frac{D_i(l_i)}{D_i'(l_i)l_i}$$

where $D_i(l_i) = \int_0^{l_i} D_i'(z)dz$ is i's residual damage function, then

PROPOSITION 3: *It is optimal for each firm i to bid its true demand curve $P_i(x_i)$ regardless of what other firms bid.*

Truth-telling is a dominant strategy for firms, so there is no need for them to form beliefs about other firms' types and/or actions. This efficient and strategy-proof result is not surprising once we realize that the auction mechanism follows a VCG payoff rule: it makes each firm i pay for its (residual) damage $D_i(l_i)$ to all other agents. This residual damage includes both the pecuniary externality imposed upon other regulated firms and the pollution externality imposed upon society.

One immediate implication of Proposition 3 is that the auction scheme implements the first-best with each firm facing the same price at the margin (i.e., $p_i = p^*$ for all i) and getting exactly the first-best allocation of licenses (i.e., $l_i = x_i^*$), which eliminates (efficiency) reasons for trading licenses after the auction. These efficiency properties can be readily seen in Figure 2: if $\hat{P}_i(x_i) = P_i(x_i)$ and $\hat{P}_{-i}(x_{-i}) = P_{-i}(x_{-i})$, then $l_i = x_i^*$, $l = x^*$ and $\hat{p} = p^*$. Although in principle the auctioneer (i.e., regulator) goes bidder after bidder determining individual prices p_i, these prices are all the same regardless of how truthful firms are (in terms of Figure 2: $p_1 = ... = p_n = \hat{p}$). But unless firms have identical demand curves, final prices, i.e., $(1-\alpha_i)p$, will differ across firms (in and off equilibrium).

III. Properties of the Mechanism

Besides being efficient and strategy-proof, the auction mechanism have additional properties that may prove relevant for purposes of practical implementation. We leave the analysis of collusion for the next section.

A. *Relationship to the VCG Mechanism*

Notwithstanding the auction mechanism follows a VCG payoff rule it is structurally different than the VCG mechanism of DHM. Without any loss of generality, let parametrize firm i's identity as $C_i(x_i) \equiv C(x_i, \theta_i)$, where θ_i is firm i's true type. In DHM firm i faces a tax schedule $T_i(x_i)$ equal to

$$T_i(x_i) = D(x_i + \sum_{j \neq i} x_j^*(\hat{\theta}_i, \hat{\boldsymbol{\theta}}_{-i})) + \sum_{j \neq i} C(x_j^*(\hat{\theta}_i, \hat{\boldsymbol{\theta}}_{-i}), \hat{\theta}_j) - A_i(\hat{\boldsymbol{\theta}}_{-i}) \tag{8}$$

where $\hat{\theta}_i$ is firm i's report to the regulator, $\hat{\boldsymbol{\theta}}_{-i}$ is the vector of firms $j \neq i$'s reports, $x_j^*(\hat{\theta}_i, \hat{\boldsymbol{\theta}}_{-i})$ is firm j's first-best pollution level as dictated by the reports of all firms, and A_i is a constant term independent of firm i's report. Although the exact value of A_i does not alter firm i's report at the margin, we know that for DHM to be a VCG mechanism the constant term A_i cannot go unspecified but must be equal to the efficient social cost had firm i not existed, that is

$$A_i(\hat{\boldsymbol{\theta}}_{-i}) = D(x_{-i}^{**}) + \sum_{j \neq i} C(x_j^{**}(\hat{\boldsymbol{\theta}}_{-i}), \hat{\theta}_j)$$

where $x_{-i}^{**} \equiv \sum_{j \neq i} x_j^{**}(\hat{\boldsymbol{\theta}}_{-i})$ and $x_j^{**}(\hat{\boldsymbol{\theta}}_{-i})$ is firm j's first-best pollution level in the absence of firm i.

Firm i's total payment under the auction mechanism and under the DHM-VCG mechanism are exactly the same in equilibrium (but not off-equilibrium), although computed in structurally different ways. This can be easily shown with the aid of Figure 2. Firm i's total payment under the auction mechanism is the shaded area to the left while the payment under the DHM mechanism is the shaded area to the right (recall that $l_i = \hat{x} - \hat{x}_{-i}$). The fact that the A_i's in DHM do not affect conduct (perhaps participation) is what distinguishes that mechanism from the one in this paper.

Despite the payoff equivalence between the two mechanisms, the structural differences will prove critical in explaining why in some other, yet important, circumstances (e.g., inelastic supply, collusive behavior) the auction mechanism continues to perform equally well, i.e., delivering the first-best, while the DHM-VCG mechanism does not.

B. *Evolution of Paybacks*

As we increase the number of firms, firm i has virtually no effect on the equilibrium price, so $D_i'(x_i^*) \approx D_i'(0)$. And by replacing $D_i''(x_i) = 0$ into (6), the differential equation's unique solution happens to be $\alpha_i(l_i) = 0$; hence, the auction scheme has converged to the Pigouvian principle for taxing externalities.

To illustrate how rapidly the auction's payment rule approaches Pigou, let us consider a numerical example. Suppose there are n symmetric firms with linear demand curves. The aggregate demand curve is $P(x) = \bar{p}(1 - x/x^0)$, where $\bar{p}$ is the choke price and x^0 is the unregulated level of pollution. The marginal damage function is $D'(x) = hx$. Solving as a function of the number of firms, we obtain

$$\alpha(n) = \frac{1}{2} \frac{\bar{p}}{(n-1)hx^0 + n\bar{p}} \tag{9}$$

If we further let the slopes of the aggregate demand and marginal damage curves be the same (i.e., $h = \bar{p}/x^0$), then eq. (9) reduces to $\alpha(n) = 1/(4n-2)$, where $n \geq 1$. The rebate for three firms is 10 percent, for ten firms is 2.6 percent, and for 100 firms is less than 0.3 percent.

C. *Perfectly Elastic/Inelastic Supply*

It is well known that if $D'(x)$ is constant a first-best policy is to charge a Pigouvian tax equal to D'. The auction mechanism is equivalent to this tax policy in that paybacks are exactly equal to zero but it still has the (practical) advantage that the socially efficient amount of pollution is set ex-ante, i.e., before pollution occurs.

More interestingly, if the supply curve is totally inelastic, say, at $x = \bar{x}$, whether because there is a genuine threshold at $\bar{x}$ or, more likely, because the regulator has no control over x, the auction mechanism still retains its truth-telling properties (note that each firm faces a elastic residual supply function).[18] This latter has two important implications. On the one hand, it makes the auction mechanism readily comparable to the (fixed-supply) private-value auctions of Vickrey (1961) and Ausubel (2004). In fact, the three auctions yield the same outcome in terms of revenues and allocations; although they are implemented in quite different ways. On the other hand, it introduces a sharp distinction between the auction mechanism and the DHM mechanism. DHM fails to yield the efficient outcome because each individual firm is no longer pivotal, i.e., its report does not affect the aggregate supply. DHM can lead, for example, to the inefficient Nash equilibrium in which each of the (no necessarily symmetric) n firms submit a null report (or that there are no clean up costs up to $\bar{x}/n$) and then emit $\bar{x}/n$. This equilibrium strategy reduces payments to zero and clearly no firm wants to deviate from it (note that above $\bar{x}$ firms face an infinitely large pollution fee).

D. *Off-Equilibrium Behavior*

Unlike in the single-firm case, in the context of multiple firms is in each firm's best interest to bid truthfully not only that portion of the demand curve around its first-best allocation x_i^* but rather a large portion of its demand curve. Even if each firm knows $D(x)$, it can no longer anticipate $l_i^* = x_i^*$ with precision because it does not know other firms' demand curves $P_{-i}(x_{-i})$ (it may be even unaware of the number of firms being regulated). To be more precise, a firm will only find it strictly optimal to bid truthfully the portion of its demand curve that is relevant for the auction clearing. Thus, if firms assign zero probability to the event that the clearing price will fall below some value, say $\underline{p}$, firms can just bid an almost perfectly elastic (or inelastic for that matter) demand curve for $p \leq \underline{p}$.[19] While this off-equilibrium behavior has no consequences on the clearing price, and hence, on implementing the first-best allocation, it does have an effect on firms' total payments. But because demand schedules are non-increasing in p, a firm's total payment will never be greater (and generally smaller) than the Pigouvian payment.

E. *Budget Balancing*

[18] There is, however, one exception: the "single" firm will no longer submit its true demand curve but $\hat{P}(x) = 0$. But since x is fixed and there is only one firm, this has no allocative implications.

[19] Going back to the fixed-supply case discussed in the preceding paragraph, the existence of $\underline{p} > 0$ could eventually break the equivalence with the Vickrey and Ausubel auctions in terms of revenues (but not in terms of allocations).

The auction mechanism is, like any other VCG mechanism, a non-budget-balanced mechanism both in and off equilibrium (unless $\hat{X}_i(p) = 0$ for all i). Although there is no efficiency reasons for balancing the budget there may be political economy reasons for doing so (Tietenberg, 2003).[20] As first pointed out by Theodore Groves and John Ledyard (1977), if there are at least three agents it should be possible to balance the budget for a variety of mechanisms. The basic idea is to distribute the surplus or deficit generated by each agent ($D_i(l_i)$ in our case) among the other agents in some lump-sum manner as to avoid any incentive effects. Behind this idea lies an implicit "separability" condition that in our case would allow us to either make the payment $D_i(l_i)$ independent of some firm j's report (i.e., $\hat{P}_j(x_j)$), as in Duggan and Roberts (2002), or to perfectly disentangle the contribution of each firm $j \neq i$'s report to firm i's payment, as in Varian (1994). By construction, the auction mechanism lacks of such separability; hence, there is no way in which the mechanism can be modified to achieve perfect budget-balancing while retaining its first-best properties.[21]

There exists, however, an approximate solution. Building upon the idea of Groves and Ledyard (1977), let denote by $D_j^{-i}(l_j^{-i})$ the total payment that firm j would have hypothetically faced under the same auction mechanism but in the absence of firm i's demand schedule, where l_j^{-i} is the corresponding number of licenses allocated to j. The regulator can thus fashion a lump-sum compensation refunding R_i for firm i using these influence-free hypothetical payments. For example,

$$R_i = \frac{1}{n-1} \sum_{j \neq i} D_j^{-i}(l_j^{-i})$$

where $n \geq 2$.

This solution assures a perfectly balanced budget (i.e., $\sum_{i=1}^{n} R_i = \sum_{i=1}^{n} D_i(l_i)$) only in the limiting case of a large number of firms; otherwise, $\sum R_i$ could be smaller, greater or equal than $\sum D_i(l_i)$. The ratio $\rho \equiv \sum R_i / \sum D_i(l_i)$ will ultimately depend on the number of firms and shape of the demands and marginal damage curves. For linear curves, for example, it can be shown that for three (symmetric) firms ρ can be anywhere between 0.60 and 1.50, for ten firms anywhere between 0.90 and 1.11 and for 100 firms anywhere between 0.99 and 1.01. Thus, a

[20] We may also want to take into consideration the general equilibrium reasons of A. Lans Bovenberg and Lawrence H. Goulder (1996) for not balancing the budget.

[21] The reason why the complete-information mechanisms of Duggan and Roberts (2002) and Varian (1994) can balance the budget is because they are based on discrete announcements by firms. In the former firms announce quantities while in the latter they announce prices. In the auction mechanism firms announce a continuum of quantity-price pairs.

regulator that cannot run a deficit, i.e., constrained to return at most $\sum D_i(l_i)$ to firms, can inform in advance that it will return only some fixed fraction of the total $\sum D_i(l_i)$ (in the case of 10 firms this fraction could be 90 percent).

F. *Shadow Cost of Public Funds*

Closely related to the previous analysis is the question of how the mechanism performs when rebates are costly to the regulator; for example, they could be used to reduce distortionary taxation somewhere else in the economy. For simplicity consider a single firm. The regulator's problem then becomes

$$\min D(l) + C(l) + \lambda\alpha(l)D'(l)l$$

where λ is the shadow cost of public funds. The full-information social optimum is implemented by setting $\alpha(l) = 0$ and the Pigouvian tax $\tau = D'(l^*) = -C'(l^*)$.

Under incomplete information, however, the regulator faces a trade-off between allocative efficiency and information rent extraction. Since the firm's problem remains unchanged (i.e., $\min C(l) + (1 - \alpha(l))pl$, where $p = D'(l)$), the payback function that best solves the regulator's trade-off is given by (its derivation follows that of Proposition 2)

$$\alpha(l) = \frac{1}{1+\lambda}\left(1 - \frac{D(l)}{D'(l)l}\right)$$

which, in turn, leads to the "shadow-cost" equilibrium condition

$$D'(l^{sc}) + C'(l^{sc}) + \frac{\lambda D''(l^{sc})l^{sc}}{1+\lambda} = 0 \tag{10}$$

Unless $D''(\cdot) = 0$, $l^{sc} < l^*$. That pollution levels are lower under asymmetric information is consistent with second-degree price discrimination principles. Note that equilibrium paybacks (i.e., information rents) are equal to $(D'(l)l - D(l))/(1 + \lambda)$, which are increasing in l, or equivalently, in $P(\cdot)$. It should be nevertheless clear that the auction mechanism is not second-best optimum because it does not pay attention to any prior information the regulator may have about the firm's type. A truly second-best scheme can at worst replicate the (incentive-compatible) auction outcome but it will generally improve upon it.[22]

G. *Dynamics and Investment*

[22] The auction mechanism still has the advantage of simplicity *vis-a-vis* a menu of a large number of contracts. Numerical exercises can shed light on welfare differences between the two schemes.

Consider two dates, $t = 1, 2$, and $n \geq 1$ firms. For notational simplicity assume no discounting. Firm i's abatement costs at date 1 are $C_i(x_i)$ but at cost I_i incurred at date 1 it can reduce its abatement costs at date 2 to $C_i(x_i, I_i)$, where I_i is the (irreversible) amount of R&D investment in more efficient technologies and $\partial C_i(x_i, I_i)/\partial I_i < 0$ for all x_i. Demand schedules are, respectively, $P_i(x_i) = -C_i'(x_i)$ and $P_i(x_i, I_i) = -\partial C_i(x_i, I_i)/\partial x_i$. The damage function in each period is $D(x)$, where $x = \sum x_i$. First-period social optimum is well known: $D'(x) + C_i'(x_i) = 0$ for all $i = 1, ..., n$. The social optimum for second-period pollution and first-period R&D is given by the first-order conditions

$$\frac{\partial C_i(x_i, I_i)}{\partial x_i} + D'(x) = 0 \tag{11}$$

$$\frac{\partial C_i(x_i, I_i)}{\partial I_i} + 1 = 0 \tag{12}$$

for all $i = 1, ..., n$.

Most regulations, whether command-and-control or market-based, will fail to yield (11)–(12) because of incomplete information (regarding abatement costs and investment) and/or time inconsistency (e.g., Gary Biglaiser et. al, 1995; Laffont and Jean Tirole, 1996). The auction mechanism is immune to both such problems. The regulator must run two separate auctions: for period-1 licenses and for period-2 licenses. Since firms do not perfectly know each other, period-2 licenses must also be auctioned off at date 1, i.e., before investments take place.[23] In allocating period-2 licenses each firm i is asked to bid a demand schedule $\hat{P}_i(x_i, I_i)$, which must be a function of the different investment levels firm i may pursue.

In deciding whether to submit its true demand schedule $P_i(x_i, I_i)$, firm i solves the problem

$$\min_{l_i, I_i} C_i(l_i(I_i), I_i) + D_i(l_i(I_i); I_{-i}) + I_i$$

where $l_i(I_i)$ is the number of licenses going to firm i, which is conditioned by its own investment, and $D_i(l_i(I_i); I_{-i})$ is the associated total payment, i.e., firm i's residual damage.[24] The first-order condition for l_i is $\partial C_i(l_i, I_i)/\partial l_i + \partial D_i(l_i)/\partial l_i = 0$ and the first-order condition for I_i

[23] If firms have complete information (but not the regulator) period-2 licenses can be auctioned off at date 2.

[24] I have deliberately included I_{-i} in $D_i(\cdot)$ to emphasize that investments by firms other than i enter optimally into $D_i(\cdot)$ as dictated by first-order conditions (11) and (12). For example, to find the equivalent of $\hat{p}_{-i}$ of Figure 2, we use the vector of $n-1$ reports $\hat{\mathbf{P}}_{-i}$ to solve (11) and (12). The rest of the curve $D_i'(\cdot)$ is constructed by computing at any $p > \hat{p}_{-i}$ the abatement/investment response of each firm $j \neq i$.

is

$$\frac{\partial C_i(l_i(I_i), I_i)}{\partial I_i} + 1 + \left(\frac{\partial C_i(l_i(I_i), I_i)}{\partial l_i} + \frac{\partial D_i(l_i(I_i))}{\partial l_i} \right) \frac{\partial l_i(I_i)}{\partial I_i} = 0$$

Since $\partial D_i(l_i)/\partial l_i = D'(l)$ for all i, the auction mechanism clearly induces firms to bid truthfully and, hence, to invest and abate optimally from a social standpoint.[25]

IV. Collusion

The workings of bidding rings or auction cartels have received a fair amount of theoretical and empirical attention in the auction literature (e.g., R. Preston McAfee and John McMillan, 1992; Klemperer, 2004). In this section I discuss how the auction scheme proposed in this paper performs under collusive behavior, if sustainable, and whether it requires of any adjustment in order to preserve its first-best properties.

The way to implement a collusive agreement in our multi-unit auction is not very different from the description of McAfee and McMillan (1992) for a single-unit auction except for some elements that I will explain below. Cartel firms need to both coordinate on their bidding schedules and agree on the procedure for sharing the cartel profits. But because cartel members do not know each other's demand curves, in implementing the collusive agreement the cartel organization must itself overcome an adverse-selection problem: it must induce its members to truthfully reveal their private information. In other words, the cartel itself faces an internal mechanism design problem. I will first present the optimal (i.e., maximal profits) collusive agreement and then an internal mechanism the cartel can use to implement such outcome.

A. *Optimal Collusive Agreement*

Imagine for a moment that there is a relatively large number of independent production plants. In the non-cooperative equilibrium each plant i operates at its first-best level x_i^* and receives virtually no payback. Imagine now that all those plants belong to a single holding company subject to the same auction scheme. The holding company is clearly better off because is not only operating at the same level (x_i^* at plant $i = 1, ..n$) but also receiving a strictly positive payback (1/2 of the auction revenues if $D'(x) = hx$). A good collusive agreement would then

[25] Note that if period-2 damages are uncertain at date1 —so the regulator operates under an expected damage function— investments will be no longer optimal ex-post, i.e., at date 2. Given those investments, the regulator can nevertheless implement optimal levels of abatement at date 2 by selling/buying additional licenses with the same auction mechanism.

be for plants to coordinate as if they were acting as a single entity. It can be established more generally, however, that

PROPOSITION 4: *The optimal collusive agreement for a cartel of $m \leq n$ firms is to submit only one serious bid with the true aggregate demand curve of the cartel, say $P_c(x_c)$. One cartel member submits the serious bid while all the other members submit empty demand schedules. The optimal collusive agreement delivers the first-best allocation.*

It is remarkable that the auction mechanism, without any sort of adjustment, can deliver the first-best even when firms are colluding. There are three interrelated reasons for that. First, paybacks for any given level of licenses are largest when the cartel faces the total supply function instead of a series of residual supply functions. Second, clean-up costs for any given level of licenses are lowest when they can be split cost-effectively across all firms in the cartel. Third, the single-firm analysis has already shown that the level of licenses that minimizes overall costs (clean-up costs and payments) is the first-best level.

These same three reasons also help explaining why cartel profits are increasing with the number of cartel members. Unlike in the single-unit auctions of McAfee and McMillan (1992), where the addition of a "low-valuation" member only contributes to dissipate cartel rents, in our multi-unit auction the most profitable cartel is an all-inclusive cartel (i.e., $m = n$). Existing members may eventually restrict additional participation insofar as it helps to prevent detection by antitrust authorities.

There are two additional observations. First, from looking at expression (8) it is not difficult to see that a collusive agreement under the DHM mechanism would depart from the first-best allocation. Because of the positive constant term $A_i(\hat{\boldsymbol{\theta}}_{-i})$ firms would coordinate on some over-reporting of their types (i.e., demand curves).[26] If we set the constant term to zero, however, collusive and non-cooperative behavior would be no different but payments would be so large that the (now Groves) mechanism would be of little practical value.

Second, there are many other, yet less profitable, collusive agreements under the auction mechanism. For example, it is possible to show that firms would benefit if they can coordinate on some demand reduction. A sub-optimal agreement like this would certainly move us away from the first best. But there is no good reason for firms to ever coordinate on a sub-optimal

[26] Note that the optimal (all-inclusive) collusive agreement under DHM would be defined by the n-report vector $\hat{\boldsymbol{\theta}}$ that solves

$$\min n \cdot \left(\sum_{i=1}^{n} C_i(x_i(\hat{\boldsymbol{\theta}}, \theta_i)) + D(x) \right) - \sum_{i=1}^{n} A_i(\hat{\boldsymbol{\theta}}_{-i})$$

where $x = \sum x_i$.

agreement if they can enforce the optimal agreement, as I argue in the next section.

B. *Implementing the Collusive Agreement*

The arguments made thus far have assumed that the cartel submits only one serious bid and this is the aggregate demand of cartel members. To do this, however, the cartel has to induce its members to truthfully reveal their individual demand curves. In addition, collusive profits have to be shared among the members in a way that the members would wish to participate in the cartel and not to deviate at the auction.

McAfee and McMillan (1992) provides a complete description of the two commonly observed forms of cartel organization: weak cartels, cartels whose members are unable to make transfer payments among themselves, and strong cartels, cartels whose members can both make transfer payments and exclude new entrants. While either type of organization may eventually arise in the single-unit auctions of McAfee and McMillan (1992), for the multi-unit auctions studied in this paper we must restrict attention to strong cartels.

It is hard to imagine how a weak cartel could achieve any degree of cooperation in the multi-unit incomplete-information environment of the auction mechanism. In the absence of after-auction transfers, a week cartel must conform itself with a sub-optimal collusive agreement in which each member (tacitly or not) agree on shading their bids to some extent. Since a firm's dominant strategy at the auction is to bid truthfully, a necessary condition for the sustainability of such an agreement is that cartel members can detect deviations at the auction. But unlike in a single-object auction, where "weak-cartel" members coordinate on bidding the seller's reserve price, detecting deviations in the multi-unit auction mechanism requires cartel members to have information on each and every member's demand curve. In the absence of transfers, it is impossible for the cartel to devise an internal (incentive-compatible) mechanism that can provide cartel members with such information prior to the auction.

Paying exclusive attention to strong cartels may seem restrictive but I would argue that in our context it is the most natural type of organization of the two because licenses are readily transferable either through bilateral contracts or in the (after-auction) spot market.[27] Thus, the transferable nature of licenses provides the cartel with the needed flexibility to structure transfer payments without necessarily facilitating its detection.[28]

[27] Although in equilibrium we should not observe any trade of licenses after the auction, in reality there will be always some trading activity as firms accomodate to factor price variation between auctions.

[28] Although collusion is less of a concern here (Proposition 4), antitrust authorities may still dislike it in that it can facilitate the spread-out of collusion through multi-market contacts.

Let us then consider a possible strong cartel of $m \leq n$ members indexed as $j = 1, ..., m$. In implementing the optimal collusive agreement of Proposition 5, the cartel organization must solve two intertwined problems. First, it must put in place an internal scheme that induces firms to truthfully reveal their individual demand curves to the cartel organization, which, as in McAfee and McMillan (1992), we will call the *cartel mechanism*. Second, the cartel must ensure obedience to the cartel mechanism, that is, it must be equipped with the ability to detect and credibly punish deviators.

Suppose for now that cartel has solved the second problem (I will come back to this shortly) and focus on the first problem. In so doing, consider the following cartel mechanism. Prior to the official auction, cartel members first agree on how to divide cartel profits by determining shares $\omega_j > 0$, where $\sum_{j=1}^{m} \omega_j = 1$. One plausible criteria can be historic use of the resource, that is $\omega_j \approx x_j^0 / \sum_{k=1}^{m} x_k^0$ (in McAfee and McMillan (1992) $\omega_j = 1/m$). Cartel profits are defined as the difference between the payment associated to the optimal agreement and the sum of the non-cooperative payments that cartel members would have faced at the auction for the same demand schedules reported to the cartel mechanism. After ω_j's are set, cartel members report their demand schedules $\tilde{P}_j(x_j)$ to the cartel mechanism. Let $\tilde{P}_c(x_c)$ denote the aggregate demand curve reported by cartel members. The cartel mechanism selects an arbitrary member to be the *serious bidder*, say bidder 1, which bids $\hat{P}_1(x_1) = \tilde{P}_c(x_1)$. Remaining cartel members bid $\hat{X}_j(p) = 0$ for all $j = 2, ..., m$.

The cartel mechanism also establishes the way licenses and payments are transferred across cartel members posterior to the auction. Let denote by l_c the number of licenses received by the serious bidder at the auction and by $D_c(l_c)$ the corresponding payment, where $D_c(x_c)$ is the residual damage function faced by the cartel. The cartel mechanism establishes that each cartel member j will receive an amount of licenses exactly equal to what he would have individually obtained at the auction for the demand curve that he reported to the mechanism. Member j's total payment for the l_j licenses will be equal to what he would have individually paid at the auction, $D_j(l_j)$, minus a fraction ω_j of the cartel profits. Payment $D_j(l_j)$ is computed as in the auction mechanism, that is, using the aggregate demand curve reported by the remaining cartel members, $\tilde{P}_{-j}(x_{-j})$. More precisely, $D_j(x_j) \equiv \int_0^{x_j} D_j'(z)dz$, where $D_j'(x_j) \equiv D_c'(x_c) - \tilde{P}_{-j}(x_{-j})$ for all j.[29]

[29] In an all-inclusive cartel, i.e., $m = n$, $D_c(x_c)$ is known prior to the auction ($D_c(x_c) = D(x)$). In a partial cartel, i.e., $m < n$, $D_c(x_c)$ is only learned posterior to the auction. Although the immediate auction results (i.e., l_c, $D_c(l_c)$ and $D_c'(l_c) = \hat{p}$) will provide the cartel with insufficient information to fully reconstruct the curve

PROPOSITION 5: *Assuming that the cartel members can agree on the* ω_j*'s, it is a dominant strategy Nash equilibrium for them to report truthfully to the cartel mechanism, i.e.,* $\tilde{P}_j(x_j) = P_j(x_j)$ *for all* $j = 1, ..., m \leq n$.

Before moving onto the cartel's second implementation problem, that of obedience with the cartel mechanism, let me briefly touch on three issues. First, the cartel mechanism proposed above is not the only (incentive-compatible) mechanism available to the cartel. It has the advantage, however, that by sharing the format of the auction mechanism it makes it easier for members to understand its workings.

Second, I have little to add on how firms will come to an agreement on the ω_j's other than pointing out that I see no reason for negotiations to fall apart because it is all about splitting spoils of unknown but positive magnitude. In other words, the bargaining process for setting the ω_j's does not involved the type of information asymmetries that are usually associated to negotiation failure (e.g., Wiggins and Libecap, 1985).

Third, the cartel mechanism (whether the one proposed here or any other) must be executed in its entirety prior to the official auction, except for the actual transfer of licenses (and payments) across cartel members. Unlike in a single-object auction where the cartel mechanism (e.g., "knockout" auction) that decides which of the cartel members will keep the object can be conducted either before or after the official auction (McAfee and McMillan, 1992), in our multi-unit environment this is simply not possible for both information and incentive reasons. On the one hand, the serious bidder must be informed of $\tilde{P}_c(x_c)$ before coming to the official auction. On the other hand, the use of any ex-post bidding procedure for determining how to allocate l_c and $D_c(l_c)$ across cartel members (e.g., a knockout auction with a structure similar to the auction mechanism but for an inelastic supply) will necessary distort member's (ex-ante) incentives in communicating with the cartel mechanism.

Let us now look at the cartel's second problem that of enforcement with the cartel mechanism. Unlike in McAfee and McMillan (1992), the cartel in our multi-unit context is coalition proof in that it requires of no patience from its members to maintain cooperation through the official auction. The optimal deviation of bidder j, whether it is the serious bidder ($j = 1$) or any of the non-serious bidders ($j = 2, ..., m$), is to bid an empty demand schedule to the cartel

$D_c(x_c)$, I see no reason why the serious bidder, or any bidder for that matter, cannot request information on $D_c(x_c)$ from the auctioneer. Furthermore, and as double-check practice, it may be a good idea to provide bidders with their own residual supply functions (in case of a mistake in the clearing of the auction, there will be always some bidder that will find it profitable to report so). If bidders are not entitled to request such information, the cartel can alternatively use a first-order (linear) approximation for $D'_c(x_c)$.

mechanism (so as to reduce its residual supply curve at the official auction to the maximum extent possible) and then bid its true demand curve $P_j(x_j)$ at the official auction. But this deviation leaves the deviating bidder strictly worse off in an amount exactly equal to its share ω_j of the otherwise cartel profits.

V. Final Remarks

I have developed an auction mechanism for the optimal regulation of a commons resource (e.g., clean air, water body, open fishery, native forest, etc.) when the regulator lacks information about the characteristics of the firms to be regulated. The mechanism is developed under the additional assumption that firms know nothing about other firms' characteristics. The mechanism is not only simple in that it is based on commonly used instruments (i.e., transferable licenses) but also remarkably effective in that it delivers the first-best allocation regardless of the number of firms and of whether they are acting noncooperatively or collusively. The mechanism yields the efficient allocation even when the aggregate supply of licenses is fixed —either because of a genuine threshold or because the regulator has no control over the aggregate supply. In addition, the mechanism provides firms with incentives to invest in socially optimal levels of R&D.

One aspect not treated in the paper is the possibility that emissions (or resource use, more generally) cannot be perfectly monitored. In addition to the adverse selection problem of not observing a firm's type (e.g., abatement costs) the regulator must now overcome the moral hazard problem of not perfectly observing the firm's action. In a recent paper, Montero (2005) compares the performance of two instruments —grandfathered transferable licenses and performance standards— in such information environment. He finds that in some cases a standard-alone policy can welfare dominate a licenses-alone policy. In many cases, though, the optimal policy is to combine licenses and standards. It would be interesting to study how the auction mechanism extends to the case of imperfect monitoring and to ask whether and to what extent it remains (second-best) optimal to auctioning off the licenses together with a minimum performance standard.

APPENDIX

PROOF OF PROPOSITION 1:

Let $1 > \omega_i > 0$ be the fraction of licenses allocated to firm i $(= 1, ..., n)$, so firm i receives an initial allocation of $\omega_i l$, where l is the total number of licenses and $\sum_{i=1}^{n} \omega_i = 1$ (as commonly observed in practice, ω_i could be proportional to historic emissions, that is $\omega_i \approx x_i^0/x^0$). The license market is assumed perfectly competitive (i.e., n large). Kwerel also requires $D''(x) > 0$. Let $\bar{s}$ be the maximum value the subsidy can take, which by construction fixes the maximum number of licenses to $\bar{l}$, where $D'(\bar{l}) = \bar{s}$. The regulator sets $\bar{s}$ sufficiently high that is always above the first-best level p^* for any possible realization of $P(x)$; otherwise, there is no point in using the scheme (in Kwerel (1977) s is unbounded but in reality we cannot let it go to infinity).

We will demonstrate that the pair $(\bar{s}, \bar{l})$ is the unique Nash-equilibrium outcome of Kwerel's scheme when licenses are grandfathered. From the arguments in the text we do not need to consider the case of under-reporting. Thus, for a reported aggregate demand curve $\hat{P}(x) \geq P(x)$, the subsidy level is s and the market price of licenses is $p = s$; hence, firm i's total compliance costs as a function of s becomes

$$TC_i(s) = C_i(X_i(s)) + s \cdot (X_i(s) - \omega_i l(s)) \tag{A1}$$

where $l(s) = D'^{-1}(s)$. The first term of (A1) is abatement cost and the second term is the net cost of purchasing licenses (which is negative when the firm is a net seller of licenses). Consider first the case in which the regulator sets $\bar{s}$ "close" to infinity. It is not difficult to see that no firm has incentives to move the outcome away from the pair $(\bar{s}, \bar{l})$. Since $X_i(s = \bar{s}) = 0$ for all $i = 1, ..n$ (firms either shut down operations or install backstop zero-emission technologies), all firms become net sellers to the government and their total costs, $TC_i(s = \bar{s}) = C_i(0) - \omega_i \bar{s}\bar{l} < 0$, reach the minimum (recall that $l'(s) > 0$). Consequently, all firms will submit infinitely large demand curves $\hat{P}_i(x_i)$ so as to ensure that $s = \bar{s}$.

Consider now the case in which $\bar{s}$ is "only large" in the sense that for some or all firms $X_i(\bar{s}) > 0$. We have that

$$\frac{dTC_i(s)}{ds} = C_i' \cdot \frac{dX_i(s)}{ds} + X_i(s) - \omega_i l(s) + s \cdot \left(\frac{dX_i(s)}{ds} - \omega_i \frac{1}{D''(l(s))} \right) \tag{A2}$$

But $C_i' = -s$, so evaluating (A2) at $s = \bar{s}$ reduces to

$$\left. \frac{dTC_i(s)}{ds} \right|_{s=\bar{s}} = X_i(\bar{s}) - \omega_i l(\bar{s}) - \frac{\omega_i \bar{s}}{D''(\bar{l})}$$

Since over-reporting leads to $\sum_{i=1}^n X_i(\bar{s}) < l(\bar{s})$, there must be a number of firms for which $dTC_i(\bar{s})/ds < 0$ (note that if firms are symmetric it is immediate that $dTC_i(\bar{s})/ds < 0$ for all firms; if firms are heterogeneous it may be still the case that $dTC_i(\bar{s})/ds < 0$ for all firms). Firms for which $dTC_i(\bar{s})/ds < 0$ have no incentives to move the outcome away from $s = \bar{s}$, hence, they will report $\hat{X}_j(p) = \infty$ for all $p \geq 0$, so as to ensure that $s = \bar{s}$. Firms for which $dTC_j(\bar{s})/ds > 0$, if any, cannot fully counterbalance these over-reporting schedules because at best they can report $\hat{X}_j(p) = 0$ for all $p \geq 0$. Therefore, the equilibrium outcome will necessarily be the pair $(\bar{s}, \bar{l})$.

PROOF OF PROPOSITION 2:

Let

$$g(l) = \exp \int \frac{D''(l)l + D'(l)}{D'(l)l} dl = \exp \int d \ln(D'(l)l) = D'(l)l$$

Then the solution to the differential equation (6) for $0 \leq l < \infty$ is given by

$$\alpha(l) = \frac{1}{g(l)} \left(K + \int g(l) \frac{D''(l)}{D'(l)} dl \right) = \frac{1}{D'(l)l} \left(K + \int D''(l) l dl \right)$$

where K is an integration constant. Integrating by parts, we obtain

$$\alpha(l) = \frac{1}{D'(l)l} (K + D'(l)l - D(l))$$

and setting the constant term K to zero finishes the proof.

PROOF OF PROPOSITION 3:

It follows immediately from the construction of $\alpha_i(l_i)$ and Proposition 2.

PROOF OF PROPOSITION 4:

Without any loss of generality let us parametrize firm i's inverse demand function as $P_i(x_i) \equiv P(x_i, \theta_i)$ where θ_i is an index of type and $\partial P/\partial \theta > 0$ (similarly, the parametrization for the demand function is $X_i(p) \equiv X(p, \theta_i)$ where $\partial X/\partial \theta > 0$). There is a one-to-one correspondence between a reported demand schedule $\hat{P}_i$ and a reported type $\hat{\theta}_i$. Consider for the moment only two firms, i and j. The firms' reports $\hat{\theta}_i$ and $\hat{\theta}_j$ conducive to the most profitable collusive agreement are found by solving

$$\min_{\hat{\theta}_i,\, \hat{\theta}_j} C(x_i, \theta_i) + C(x_j, \theta_j) + [1 - \alpha_i(l_i(\hat{\theta}_i, \hat{\theta}_j))] \hat{p}(\hat{\theta}_i, \hat{\theta}_j) l_i(\hat{\theta}_i, \hat{\theta}_j) + [1 - \alpha_j(l_j(\hat{\theta}_i, \hat{\theta}_j))] \hat{p}(\hat{\theta}_i, \hat{\theta}_j) l_j(\hat{\theta}_i, \hat{\theta}_j)) \tag{A3}$$

subject to

$$x_i + x_j = l_i(\hat{\theta}_i, \hat{\theta}_j) + l_j(\hat{\theta}_i, \hat{\theta}_j) = l(\hat{\theta}_i, \hat{\theta}_j) \tag{A4}$$

where $\hat{p}(\hat{\theta}_i, \hat{\theta}_j) \equiv \hat{p}$ is the auction clearing price as a function of firms' bids and $l_i(\hat{\theta}_i, \hat{\theta}_j) \equiv l_i$ is the number of licenses allocated to firm i. In what follows I will omit $\hat{\theta}_i$ and $\hat{\theta}_j$ unless it would otherwise cause confusion. From Proposition 2 we know that

$$[1 - \alpha_i(l_i)]\hat{p}l_i = D_i(l_i) = \hat{p}l_i - \int\limits_{\hat{p}_j(\hat{\theta}_j)}^{\hat{p}} [X^s(p) - X(p, \hat{\theta}_j)]dp \tag{A5}$$

where $X^s(p)$ is the social supply function, i.e., $D'(x)$, so $X^s(p) - X(p, \hat{\theta}_j)$ is the residual supply faced by firm i, i.e., $D_i'(x_i, \hat{\theta}_j)$; and $\hat{p}_j(\hat{\theta}_j) \leq \hat{p}$ is the (hypothetical) clearing price in the absence of firm i's bid (in terms of Figure 2, $\hat{p}_j(\hat{\theta}_j)$ corresponds to $\hat{p}_{-i}$). The first-order condition for (A3) is (allowing for corner solutions)

$$\frac{\partial C(x_i, \theta_i)}{\partial x_i}\frac{dx_i}{dl}\frac{\partial l}{\partial \hat{\theta}_i} + \frac{\partial C(x_j, \theta_j)}{\partial x_i}\frac{dx_j}{dl}\frac{\partial l}{\partial \hat{\theta}_i} + \frac{\partial D_i(l_i)}{\partial \hat{\theta}_i} + \frac{\partial D_j(l_j)}{\partial \hat{\theta}_i} \geq 0 \tag{A6}$$

Recall that $-\partial C(x_i, \theta_i)/\partial x_i = P(x_i, \theta_i)$. To obtain an expression for dx_i/dl use (A4) and note that collusion optima requires

$$P(x_i, \theta_i) = P(x_j = l - x_i, \theta_j) = P(l, \theta_{i+j}) \tag{A7}$$

where $P(l, \theta_{i+j})$ is the true aggregate demand function. Totally differentiating (A7) with respect to l and rearranging leads to

$$\frac{dx_i}{dl} = \frac{P_j'}{P_i' + P_j'} \tag{A8}$$

where $P_i' \equiv \partial P(x_i, \theta_i)/\partial x_i$. To obtain expressions for $\partial D_i(l_i)/\partial \hat{\theta}_i$ and $\partial D_j(l_j)/\partial \hat{\theta}_i$ (or $\partial D_i(l_i)/\partial \hat{\theta}_j$), on the other hand, note that from (A5) we have

$$\frac{\partial D_i(l_i)}{\partial \hat{\theta}_i} = \frac{\partial \hat{p}}{\partial \hat{\theta}_i}l_i + \hat{p}\frac{\partial l_i}{\hat{\theta}_i} - \frac{\partial \hat{p}}{\partial \hat{\theta}_i}[X^s(\hat{p}) - X(\hat{p}, \hat{\theta}_j)] \tag{A9}$$

But $X^s(\hat{p}) - X(\hat{p}, \hat{\theta}_j) = l_i$, so (A9) reduces to

$$\frac{\partial D_i(l_i)}{\partial \hat{\theta}_i} = \hat{p}\frac{\partial l_i}{\hat{\theta}_i} \tag{A10}$$

Similarly,

$$\frac{\partial D_i(l_i)}{\partial \hat{\theta}_j} = \frac{\partial \hat{p}}{\partial \hat{\theta}_j}l_i + \hat{p}\frac{\partial l_i}{\hat{\theta}_j} - \frac{\partial \hat{p}}{\partial \hat{\theta}_j}l_i + \frac{\partial \hat{p}_j(\hat{\theta}_j)}{\partial \hat{\theta}_j} \cdot 0 + \int\limits_{\hat{p}_j}^{\hat{p}} \frac{\partial X(p, \hat{\theta}_j)}{\partial \hat{\theta}_j}dp$$

Rearranging and inverting i by j leads to

$$\frac{\partial D_j(l_j)}{\partial \hat{\theta}_i} = \hat{p}\frac{\partial l_j}{\partial \hat{\theta}_i} + \int_{\hat{p}_i}^{\hat{p}} \frac{\partial X(p, \hat{\theta}_i)}{\partial \hat{\theta}_i} dp \tag{A11}$$

Note that since $\hat{p}_i \leq \hat{p}$ the last term of (A11) is non-negative. Plugging (A8), (A10) and (A11) into (A6), using (A7) and rearranging, the two first-order conditions become

$$\frac{\partial l(\hat{\theta}_i, \hat{\theta}_j)}{\partial \hat{\theta}_i}[-P(l(\hat{\theta}_i, \hat{\theta}_j), \theta_{i+j}) + \hat{p}(\hat{\theta}_i, \hat{\theta}_j)] + \int_{\hat{p}_i(\hat{\theta}_i)}^{\hat{p}(\hat{\theta}_i, \hat{\theta}_j)} \frac{\partial X(p, \hat{\theta}_i)}{\partial \hat{\theta}_i} dp \geq 0 \quad \text{for } i \text{ and } j \tag{A12}$$

By inspection of (A12) one arrives at two possible solutions. One solution is for i to report $\hat{\theta}_i = \theta_{i+j}$ (i.e., the true aggregate demand curve) and for j to report the corner $\hat{\theta}_j = \emptyset$ (i.e., $\hat{X}_j = 0$ for all p). If so, $P(l, \theta_{i+j}) = \hat{p} = \hat{p}_i > \hat{p}_j = D'(0)$, and hence, the first-order condition for i equals to zero and the first-order condition for j is strictly positive. The second solution is just the inverse. Both solutions are equally optimal (for the firms) and, more importantly, they implement the first-best in that firms find it in their best collusive interest to submit the aggregate true curve. Extending the proof to the case of more than two firms and to the possibility of partial collusion (i.e., collusion among a subset of firms) is straightforward.

PROOF OF PROPOSITION 5:

Given members' obedience to the cartel mechanism for whatever demand schedules they choose to report, cartel member $j = 1, ..., m \leq n$ will report the demand schedule $\tilde{P}_j(\cdot)$ that solves (recall the one-to-one correspondence between reporting $\tilde{P}_j(\cdot)$ and requesting l_j licenses)

$$\min_{l_j} C_j(l_j) + D_j(l_j) - \omega_j \left(\sum_{k=1}^{m} D_k(l_k(l_j)) - D_c(l_c) \right)$$

where $l_k(l_j)$ is member k's license allocation as a function of j's allocation and $l_c = \sum_{k=1}^{m} l_k$. The first-order condition is

$$C_j'(l_j) + D_j'(l_j) - \omega_j \left(\sum_{k=1}^{m} (D_k'(l_k) - D_c'(l_c)) \frac{dl_k}{dl_j} \right) = 0$$

But from the auction clearing condition we have $\hat{p} = D_c'(l_c) = D_k'(l_k)$ for all $k = 1, ..., m$, which finishes the proof.

REFERENCES

Ausubel, Lawrence M. "An Efficient Ascending-Bid Auction for Multiple Objects." *American Economic Review*, 2004, *94*(5), pp. 1452-75.

Biglaiser, Gary; Horowitz, John K. and Quiggin, John. "Dynamic Pollution Regulation." *Journal of Regulatory Economics*, 1995, *8*(1), pp. 33-44.

Bovenberg, A. Lans and Goulder, Lawrence H. "Optimal Environmental Taxation in the Presence of Other Taxes: General Equilibrium Analysis." *American Economic Review*, 1996, *86*(4), pp. 985-1000.

Caffera, Marcelo and Dubra, Juan. "Getting Polluters to Tell the Truth." Universidad de Montevideo, Department of Economics, Working Paper, October 2006.

Cramton, Peter and Kerr, Suzi. "Tradeable Carbon Permit Auctions: How and why to Auction not Grandfather." *Energy Policy*, 2002, *30*(4), pp. 333-45.

Dasgupta, Partha; Hammond, Peter and Maskin, Eric. "On Imperfect Information and Optimal Pollution Control." *Review of Economic Studies*, 1980, *47*(5), pp. 857-60.

Duggan, John and Roberts, Joanne. "Implementing the Efficient Allocation of Pollution." *American Economic Review*, 2002, *92*(4), pp. 1070-78.

Gersbach, Hans and Requate, Till. "Emission Taxes and Optimal Refunding Schemes." *Journal of Public Economics*, 2004, *88*(3-4), pp. 713-25.

Green, Jerry R. and Laffont, Jean-Jacques. *Incentives in public decision making.* Amsterdam: North-Holland, 1979.

Groves, Theodore and Ledyard, John. "Optimal Allocations of Public Goods: A Solution to the "Free Rider" Problem." *Econometrica*, 1977, *45*(4), pp. 783-809.

Joskow, Paul L.; Schmalensee, Richard and Bailey, Elizabeth M. "The Market for Sulfur Dioxide Emissions." *American Economic Review*, 1998, *88*(4), pp. 669-85.

Kim, Jae-Cheol and Chang, Ki-Bok. "An Optimal Tax/Subsidy for Output and Pollution Control under Asymmetric Information in Oligopoly Markets." *Journal of Regulatory Economics*, 1993, *5*(2), pp. 183-97.

Klemperer, Paul. *Auctions: Theory and practice.* Princeton: Princeton University Press, 2004.

Kremer, Ilan and Nyborg, Kjell G. "Divisible-Good Auctions: The Role of Allocation Rules." *RAND Journal of Economics*, 2004, *35*(1), pp. 147-59.

Kwerel, Evan. "To Tell the Truth: Imperfect Information and Optimal Pollution Control." *Review of Economic Studies*, 1977, *44*(3), pp. 595-601.

Laffont, Jean-Jacques and Tirole, Jean. "Pollution Permits and Environmental Innovation." *Journal of Public Economics*, 1996, *62*(1-2), pp. 127-40.

Mas-Colell, Andreu; Whinston, Michael D. and Green, Jerry R. *Microeconomics theory.* Oxford: Oxford University Press, 1995.

McAdams, David. "Adjustable Supply in Uniform Price Auctions: The Value of Non-Commitment." MIT, Sloan School of Management, Working Paper No. 4522-05, March 2005.

McAfee, R. Preston and McMillan, John. "Bidding Rings." *American Economic Review*, 1992, *82*(3), pp. 579-99.

Milgrom, Paul. *Putting auction theory to work.* Cambridge: Cambridge University Press, 2004.

Montero, Juan-Pablo. "Pollution Markets with Imperfectly Observed Emissions." *RAND Journal of Economics*, 2005, *36*(3), pp. 645-60.

Montero, Juan-Pablo. "A Simple Auction Mechanism for the Optimal Allocation of the Commons." MIT, CEEPR, Working Paper No. 2006-08, May 2006.

Roberts, Marc J. and Spence; Michael. "Effluent Charges and Licenses under Uncertainty." *Journal of Public Economics*, 1976, *5*(3-4), pp. 193-208.

Spulber, Daniel F. "Optimal Environmental Regulation under Asymmetric Information." *Journal of Public Economics*, 1988, *35*(2), pp. 163-81.

Tietenberg, Tom. "The Tradable-Permits Approach to Protecting the Commons: Lessons for Climate Change." *Oxford Review of Economics Policy* , 2003, *19*(3), 400-419.

Varian, Hal R. "A Solution to the Problem of Externalities when Agents are Well-Informed." *American Economic Review*, 1994, *84*(5), pp. 1278-93.

Vickrey, William. "Counterspeculation, Auctions and Competitive Sealed Tenders." *Journal of Finance*, 1961, *16*(1), pp. 8-37.

Weitzman, Martin L. "Prices vs. Quantities." *Review of Economic Studies*, 1974, *41*(4), pp. 477-91.

Wiggins, Steven N. and Libecap, Gary D. "Oil Field Unitization: Contractual Failure in the Presence of Imperfect Information." *American Economic Review*, 1985, *75*(3), pp. 368-85.

Wilson, Robert. "Auctions of Shares." *Quarterly Journal of Economics*, 1979, *93*(4), pp. 675-89.

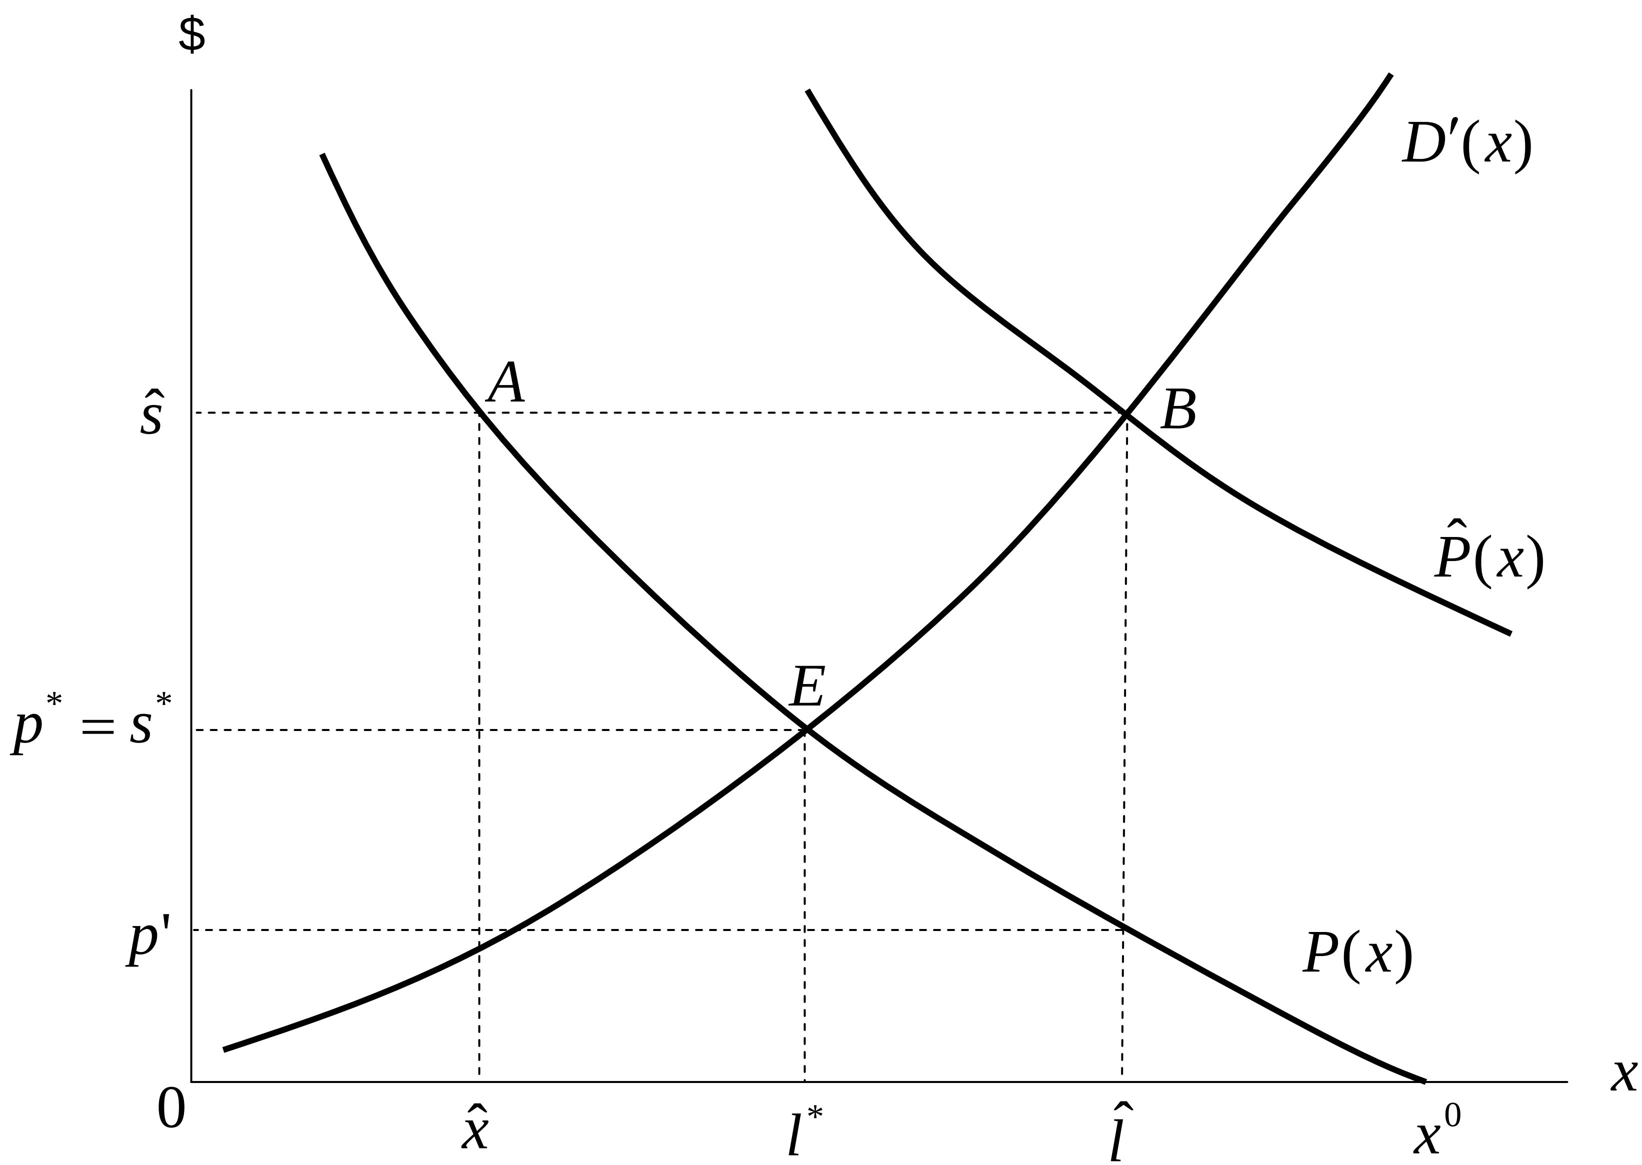

Figure 1a: Incentives to over-report

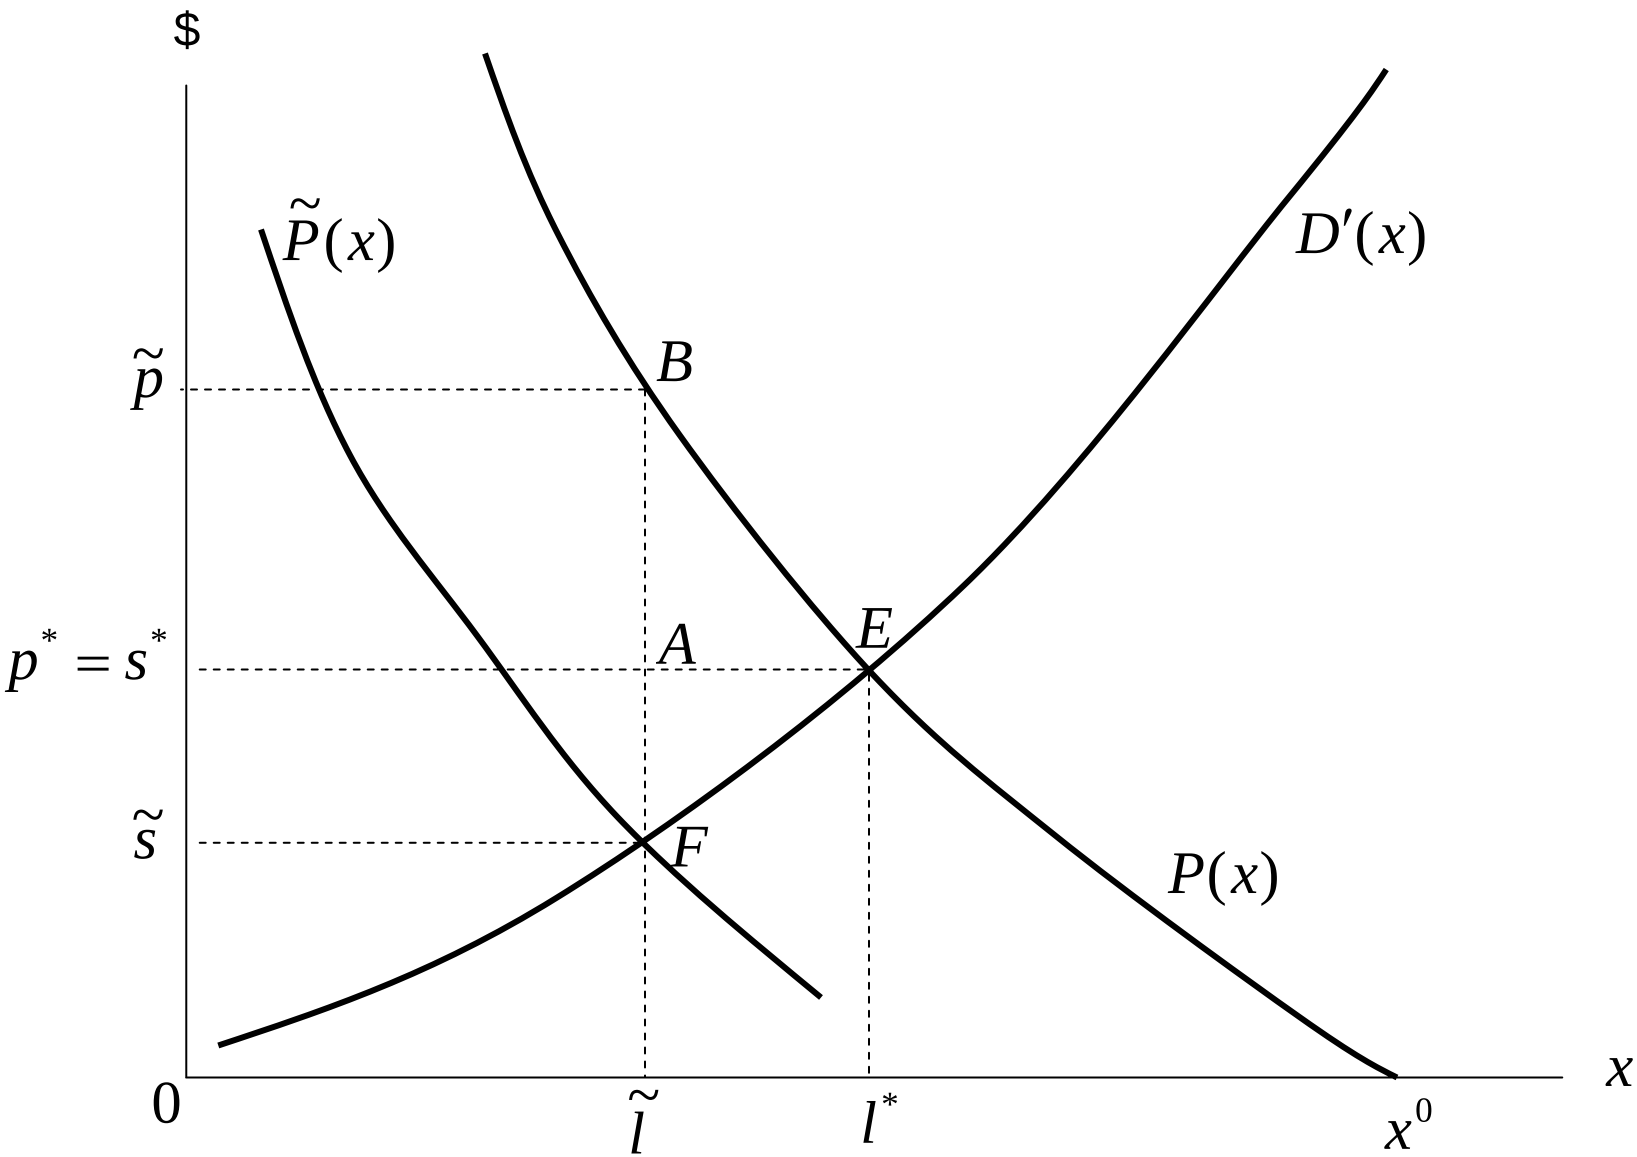

Figure 1b: Incentives to under-report

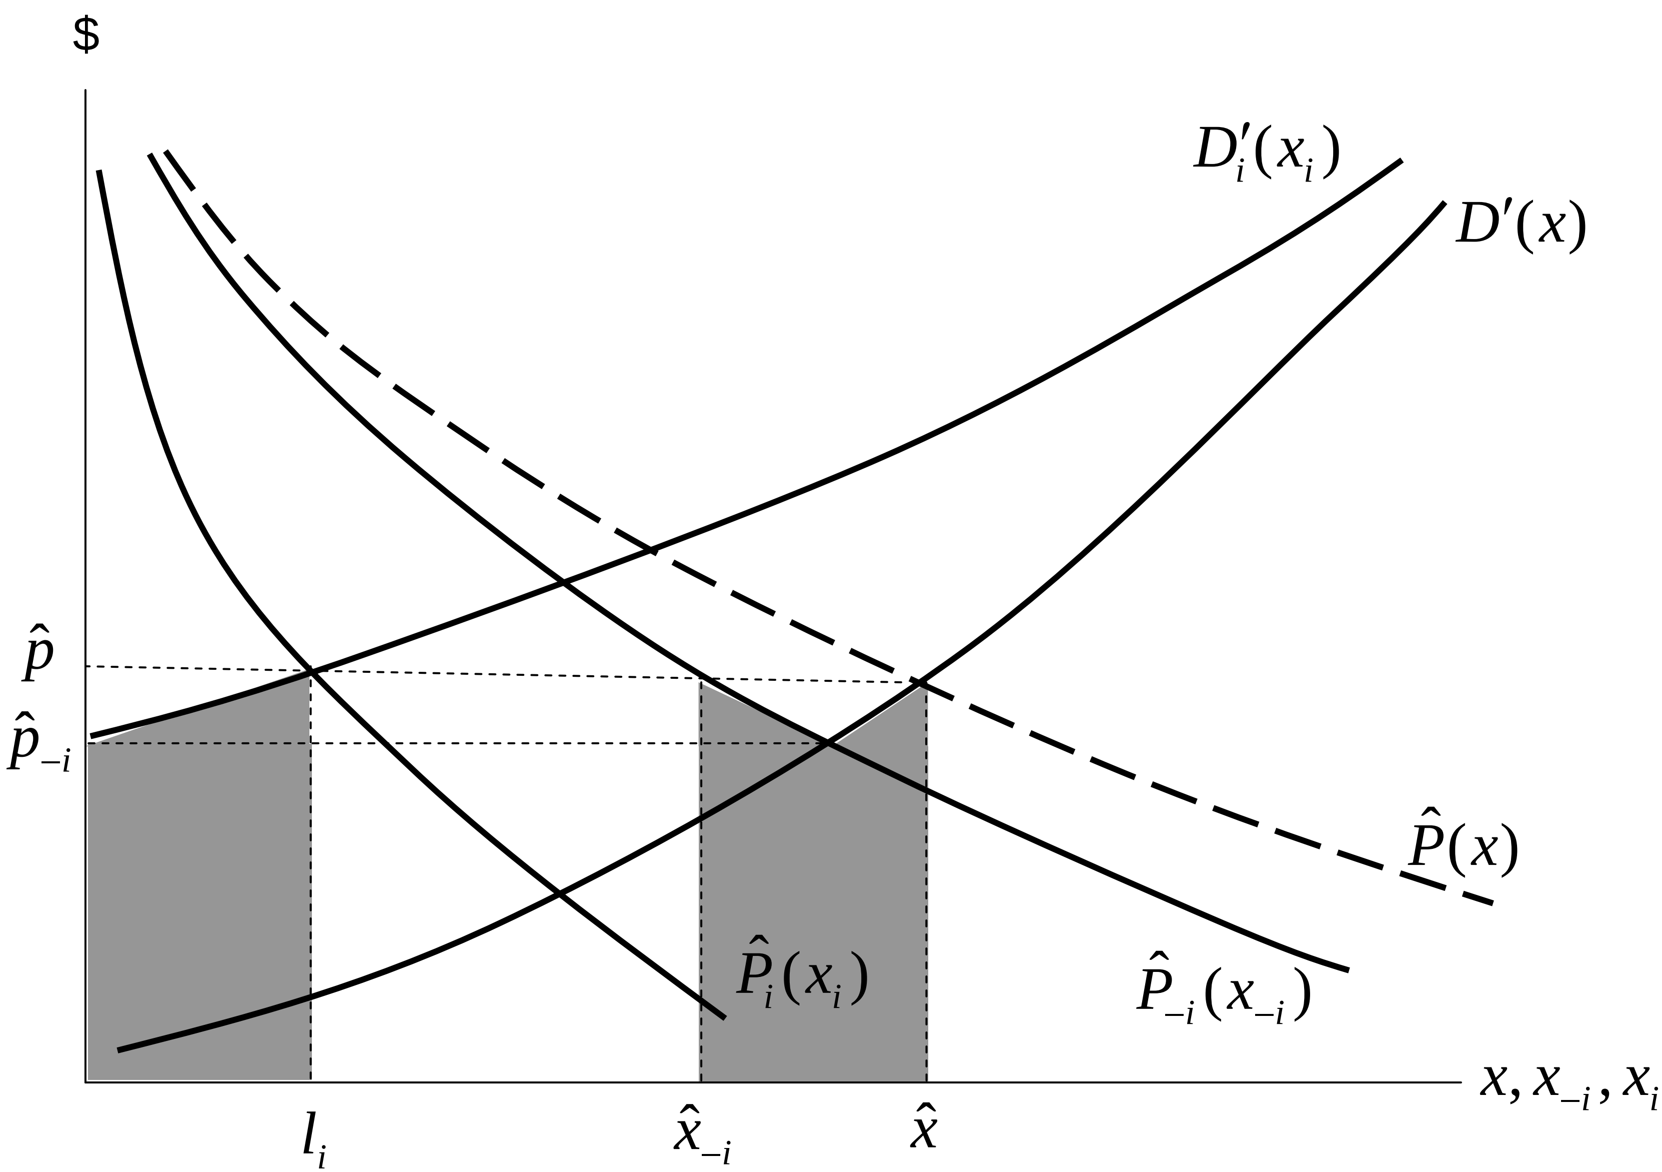

Figure 2: Residual supply (i.e., marginal damage) function