

SIMPLE PRICING SCHEMES FOR POLLUTION CONTROL UNDER ASYMMETRIC INFORMATION

Peter Kennedy

University of Victoria

Benoit Laplante

World Bank

Dale Whittington

University of North Carolina

7 January 2010

ABSTRACT

The theoretical literature on environmental regulation has proposed a variety of innovative and elegant schemes for dealing with asymmetric information, yet these schemes are rarely used in practice. This is likely due to their complexity. In contrast, relatively simple pricing schemes based on uniform increasing marginal price (UIMP) schedules are common. These schemes often take the form of a uniform discrete block pricing (DBP) schedule. We show that such schemes can be motivated as a relatively simple – albeit second-best – solution to the asymmetric information problem. We characterize the optimal linear UIMP scheme and the optimal DBP scheme, and characterize their relationship to the Pigouvian tax. We show that the optimal linear UIMP schedule is a limiting case of the optimal DBP scheme. We argue that such schemes are likely to have much greater practical appeal to regulators than more sophisticated but complex pricing schemes proposed in the existing theoretical literature.

1. INTRODUCTION

The literature on pollution pricing under asymmetric information has proposed a variety of innovative regulatory schemes derived from mechanism design theory, but these schemes are rarely used in practice.¹ A key obstacle to the implementation of these policies is their complexity, as perceived by the regulators tasked with drafting actual statutes, and by regulated firms who demand regulatory simplicity. In particular, these schemes typically require the implementation of firm-specific pricing whereby each regulated entity faces a pricing scheme tailored to fit its own individual characteristics. This requirement is not easily reconciled with the practical realities of real-world regulation.

Many regulators have nonetheless shown a willingness to use pricing schemes that go beyond a single unit price for pollution. A number of jurisdictions within the U.S. and in other parts of the world use “stepped rates” – or discrete block pricing (DBP) schemes – whereby a given unit price for pollution is applied up to a firm-specific level for each polluter, and a higher unit price is applied to any pollution from that polluter above the specified level. Similar pricing schemes are often used by water and electricity utilities.² These schemes typically apply the same price schedule to all firms and so retain a degree of simplicity sufficient to allow manageable implementation. The most common rationale for such schemes is to create strong abatement incentives for large polluters while not imposing a high unit price on all polluters.

In this paper we show that a DBP-type scheme can also provide an imperfect but relatively simple approach to regulation under asymmetric information, one that does not require a dramatic departure from actual regulatory practice. The pricing schemes we develop allow the regulator to apply the same price schedule to all firms – a property we will call *uniform treatment* – and hence possess a simplicity of structure that might facilitate actual adoption. The obvious downside with these schemes is that different sized polluters pay different prices on their marginal unit of pollution. Consequently, marginal abatement costs (MACs) are not equated across pollution sources, and marginal

¹ See Lewis (1996) for a comprehensive discussion of these regulatory schemes.

² Kaplow and Shavell (2002) point out that such schemes are also familiar to regulators in the context of progressive income tax systems.

damage may be higher or lower than those MACs. The policy design problem is to balance the costs of this shortcoming with the information-related benefits of a price schedule in which the marginal price of pollution rises with each firm's level of pollution. This policy design problem is the focus of our paper.

We address this problem in the context of an admittedly restrictive model in which MACs and marginal damage are linear, and in which the information asymmetry is limited to the intercept of the MAC schedule (rather than the slope). These limitations necessarily restrict the generality of our specific results, but they allow the derivation of sharp analytical solutions which in turn shed light on the key elements of the policy design problem.

We study two types of pricing schemes. The first is a true DBP scheme with a finite number of price steps. We derive an analytical solution for the optimal DBP scheme for any given number of price steps, and then show that under particular conditions, the optimal number of price steps is infinite; that is, the optimal scheme is a continuous, and linear, increasing marginal price (LIMP) schedule. We then examine this LIMP scheme and show that its optimal structure strikes a balance between the information-related benefits of increasing marginal prices on one hand, and an increase in aggregate abatement cost on the other, due to the non-equalization of MACs across firms.

We then relate our price schemes to those proposed by Weitzman (1978) and Roberts and Spence (1976). Weitzman examines a hybrid price-quantity scheme in which a traditional Pigouvian tax is combined with a penalty for deviating from a prescribed firm-specific quantity target. The hybrid scheme effectively provides a "safety valve" for a quantity target that is too restrictive *ex post*, and at the same time limits the overshooting that would arise from a tax that is too lax *ex post*. We show that the Weitzman scheme is equivalent to a LIMP scheme in which each regulated firm faces a different LIMP schedule, with an intercept tailored to its individual expected abatement cost. This discrimination across firms means that expected MACs are equated across sources, and hence that expected abatement cost is minimized. This is a key advantage over a LIMP scheme with uniform treatment like the one we examine, but it comes at the cost of increased regulatory complexity; the regulator cannot simply specify a price schedule that applies to all firms.

Roberts and Spence (1976) propose a different sort of hybrid price-quantity scheme, using a combination of charges and tradable licenses. Their scheme can be interpreted as a DBP scheme in which firms are entitled to trade unused portions of the price blocks. Adding the possibility of trade improves the performance of a DBP scheme because trade ensures that MACs are equated across firms in equilibrium, and this equality holds *ex post*, not just in expectation. On the other hand, trade adds an element of administrative complexity to the pricing scheme that regulators may find unattractive relative to a DBP scheme without trade. Moreover, the optimal policy parameters in the Roberts and Spence scheme can be calculated analytically only under very special conditions.

The rest of our paper proceeds as follows. In section 2 we present the model in which we examine the regulatory problem. In Section 3 we characterize the Pigouvian tax in the context of that model as a useful benchmark for the analysis of the DBP-type schemes. In section 4 we derive the optimal DBP scheme and its limiting form as a LIMP schedule. In section 5 we derive the optimal LIMP schedule under more general conditions. In Section 6 we relate our pricing schemes to those proposed by Roberts and Spence (1976) and Weitzman (1978), and in section 7 we examine the relative performance of these schemes. Section 8 provides some concluding remarks. All reported proofs are contained in the Appendix.

2. THE MODEL

Our model has N regulated firms. The MAC for firm i is given by

$$(1) \quad MAC_i(e_i) = \frac{x_i - e_i}{a}$$

where e_i is emissions by firm i , x_i is the unregulated (or no-abatement) level of emissions for that firm, and a is a positive parameter. Thus, $x_i - e_i$ measures abatement. Note that firms differ with respect to their MAC intercepts (x_i) but share the same slope

parameter.³ We assume that individual firms know their own x_i but the regulator faces uncertainty over this parameter. In particular, from the perspective of the regulator,

$$(2) \quad x_i = z_i + v_i + \eta$$

where z_i is observable, v_i is an unobserved idiosyncratic random variable, and η is an unobserved industry-wide random variable. We assume that η is drawn from a distribution with $E[\eta] = 0$ and $E[\eta^2] = \rho^2$; the v_i 's are drawn from independent and identical distributions with $E[v_i] = 0 \forall i$ and $E[v_i^2] = \omega^2 \forall i$; and η and v_i are independent. It is worth noting that these assumptions on the random variables together mean that $E[x_i] = z_i$. We let \bar{v} denote the mean of the *realized* v_i 's. Note that \bar{v} will generally *not* be equal to zero (if N is finite and $\omega^2 > 0$); it is a random variable from the perspective of the regulator.⁴ This has important implications for the policy design problem under the Roberts and Spence (1976) scheme in particular.

Firms differ according to the realized value of their v_i parameter (which we call *unobservable heterogeneity*) but they also potentially differ according to their observable parameter z_i (which we call *observable heterogeneity*). We will see later that this distinction is important for the policy design problem. The mean and variance of the observed z_i 's are denoted $z > 0$ and σ^2 respectively.

We assume the following proportional marginal environmental damage schedule:

$$(3) \quad MD = \delta \sum_{i=1}^N e_i$$

³ This specification of marginal abatement cost is common in the literature; for example, see Adar and Griffin (1976), Blair (1985) and Stavins (1996). Our assumption that the slope of the MAC function is the same across firms and known by the regulator is clearly restrictive. We adopt it because it facilitates the derivation of analytical solutions. Relaxing this assumption means that our optimal pricing schemes must be derived numerically, as do the alternative schemes to which we compare our results in sections 6 and 7. This is not necessarily an obstacle to their implementation, but it does render comparative analysis less transparent. For treatments of optimal regulation under slope uncertainty (though not in the context of DBP-type schemes), see Watson and Ridker (1984) and Hoel and Karp (2001).

⁴ One should think of \bar{v} as the mean of a sample of size N drawn from a population whose mean is zero and whose variance is ω^2 .

where $\delta \geq 0$. Emissions are assumed to be observable and verifiable.⁵ Throughout our analysis we assume that the first-best solution is always an interior one (with positive emissions for all firms).⁶

3. THE PIGOUVIAN TAX

The Pigouvian tax – levied on a per unit basis – is the textbook benchmark for pollution pricing, and it is useful to characterize the properties of the tax in this setting in order to later compare it to the other pricing schemes. It should be noted that the Pigouvian tax we derive here is the *second-best* Pigouvian tax, as distinct from the first-best tax that would be levied in the absence of uncertainty about abatement cost, because it is set *ex ante* (before the uncertainty is resolved). For the sake of brevity, we henceforth simply use the term “Pigouvian tax” to refer to this second-best per unit tax. Moreover, we restrict attention to a setting in which the regulator makes a once-and-for-all choice of the tax rate (rather than one in which the rate can be adjusted over time, a possibility discussed further below).

If the regulator imposes a per unit tax t on emissions, then the cost-minimizing response by firm i is

$$(4) \quad e_i(t) = x_i - at$$

The regulator’s problem is to choose the tax rate that minimizes expected social cost (aggregate abatement cost plus environmental damage), given the response behaviour in equation (4):

$$(5) \quad \min_t E[C(t) + D(t)] \quad s.t. \quad e_i(t) = x_i - at \quad \forall i$$

where $C(t)$ is aggregate abatement cost, given by

⁵ If emissions are observable, then it might seem reasonable to suppose that the “no abatement” emissions level should also be observable prior to the implementation of policy. In reality, most firms currently face some form of regulation that causes actual emissions to deviate from their no-abatement emissions. We envisage a setting where the regulator is moving to a new form of regulation based on emissions pricing, and where current emissions may not correspond to the no-abatement level, due to existing regulations.

⁶ This is not an unrestrictive assumption. In particular, it *must* be violated when $\omega^2 > 0$ or $\sigma^2 > 0$ and N is sufficiently large (because our MAC has a finite intercept). This means that one cannot infer anything from taking the limit as $N \rightarrow \infty$ in any of our results. This limitation is not especially important from a practical perspective, but it does explain why the reader might arrive at puzzling results if tempted to examine this limiting case.

$$(6) \quad C(t) = \sum_{i=1}^N \left[\frac{[x_i - e_i(t)]^2}{2a} \right]$$

and $D(t)$ is environmental damage, given by

$$(7) \quad D(t) = \frac{\delta \left[\sum_{i=1}^N e_i(t) \right]^2}{2}$$

Given our assumptions on the random variables, the solution to this problem is

$$(8) \quad t^* = \frac{z\delta N}{1 + a\delta N}$$

In response to this tax, emissions from firm i are

$$(9) \quad e_i(t^*) = \frac{x_i + (x_i - z)a\delta N}{1 + a\delta N}$$

In comparison, first-best (full information) emissions from firm i are

$$(10) \quad e_i^* = \frac{x_i + (x_i - x)a\delta N}{1 + a\delta N}$$

where $x = \sum_{i=1}^N x_i / N$. Thus, emissions from a given firm under the Pigouvian tax will differ from the first-best emissions level for that firm by an amount proportional to the difference between the average of the true no-abatement emission levels for the industry (x) and the regulator's expectation of that industry average (z).

The potential for error under the Pigouvian tax is illustrated in Figure 1, where emissions are measured on the horizontal axis. It illustrates a simple case with only one regulated firm and just two possible states of the world. In state A the firm has MAC_A and in state B the firm has MAC_B . The Pigouvian tax is based on the expected MAC, labeled $E[MAC]$. The firm will respond to that tax with emissions $e_A(t^*)$ if it has MAC_A and with emissions $e_B(t^*)$ if it has MAC_B . In comparison, the first-best emission levels are e_A^* and e_B^* respectively. Thus, the expected welfare loss under the Pigouvian tax – relative to first-best – is the probability-weighted sum of the two shaded areas.

Figure 1 illustrates the basic problem the regulator faces in attempting to price emissions under uncertainty, about abatement costs; it also serves to illustrate a potential solution. If the regulator abandons the per-unit Pigouvian tax approach and instead

simply presents the firm with a price schedule that corresponds to the MD function then it is clear from the figure that the firm will always choose the first-best level of emissions, regardless of whether its true MAC schedule is MAC_A or MAC_B .

This well-known result – identified by Dasgupta *et. al.* (1980) and expounded more accessibly by Kaplow and Shavell (2002) – illustrates the information benefits of using a pricing scheme in which the marginal price on emissions rises with the level of emissions. If the firm in Figure 1 has MAC_A then a lower marginal price is needed to ensure the right balance between MAC and MD; conversely, if the firm has MAC_B then a higher price is required. A price schedule in which the marginal price rises with the level of emissions allows this lower or higher price to arise endogenously without the regulator needing to know the true MAC schedule. In the case of a single firm, this kind of pricing scheme achieves the first-best solution.

The pricing problem becomes more complicated when the setting involves two or more firms, because MACs will generally not be equated across firms – and so aggregate abatement cost will not be minimized – unless all firms face the same price on their marginal unit of emissions. This condition cannot be satisfied if different firms face the same increasing marginal price scheme. One potential solution to this problem is to use a hybrid tax-quantity scheme that effectively confronts different polluters with a single price [Roberts and Spence (1976) and Weitzman (1978)]. These schemes add a layer of regulatory complexity beyond simple pricing schedules with uniform treatment. An alternative solution is to set a single Pigouvian tax rate for all firms but to adjust that tax rate over time towards its optimal level in response to the observed behaviour of firms [Karp and Livernois (1994) and Kaplow and Shavell (2002)]. This approach is vulnerable to strategic gaming by the regulated firms and requires continual regulatory revisions. Our goal in the following sections is to design a once-and-for-all pricing scheme that retains the regulatory simplicity of a policy in which all firms in the industry face the same fixed price schedule but which optimally balances the information benefits of a rising marginal price with the costs associated with forcing MACs to differ across firms in equilibrium.

4. DISCRETE BLOCK PRICING

We begin with a DBP scheme with a finite number of blocks (or steps) and associated unit prices. (Most such schemes in practice have just two blocks). Critically, the same price schedule is applied to all firms in the regulated industry or region. We will refer to such a scheme as exhibiting *uniform treatment* (a UDBP scheme).

To enable the derivation of an explicit analytical solution for the optimal UDBP scheme, we initially restrict attention to a simplified case in which $\omega^2 > 0$, but $\sigma^2 = 0$ and $\rho^2 = 0$. This implies unobservable heterogeneity across firms ($\omega^2 > 0$), where all firms are the same in expectation ($\sigma^2 = 0$), and no industry-wide uncertainty ($\rho^2 = 0$). This is the simplest possible setting in which to capture both uncertainty and heterogeneity across firms, and thus create a meaningful trade-off in the policy design problem. These assumptions are retained for the remainder of section 4.

Consider a UDBP scheme with m discrete price blocks and m associated unit prices, of the form:

$$(11) \quad p_1 \text{ for } e_i \in [0, b_1] \quad \text{and} \quad p_j \text{ for } e_i \in (b_{j-1}, b_j] \text{ for } j = 2, \dots, m$$

That is, firm i pays a price p_1 per unit on its emissions up to b_1 , a higher unit price p_2 on its emissions greater than b_1 but not exceeding b_2 , a still higher unit price p_3 on its emissions greater than b_2 but not exceeding b_3 , and so on. Such a scheme is illustrated as the step function labeled $UDBP^*$ in Figure 2, for the case of $m = 3$; ignore the line labeled $p^*(e)$ for now. Facing such a price scheme, firm i will respond in the following way:

$$(12) \quad e_i(p, b) = \begin{cases} x_i - ap_j & \text{if } [b_{j-1} + ap_j < x_i < b_j + ap_j \\ b_j & \text{if } [b_j + ap_j \leq x_i \leq b_j + ap_{j+1} \end{cases}$$

The regulator's problem is to choose $\{p_j, b_j\}$ to minimize the sum of $E[D(p, b)]$ and $E[C(p, b)]$. The first-order conditions for this problem are given by equations (A2) – (A6) in the Appendix. It is generally not possible to solve these conditions for a closed form solution, even for a specified distribution for v_i . An exception is where v_i has a uniform distribution, and the following proposition describes the optimal UDBP scheme in that case.

PROPOSITION 1. If $\sigma^2 = 0$ and $\rho^2 = 0$, and v_i is distributed uniformly with support $[-\xi, \xi]$, then the optimal UDBP scheme is given by $\{p_j^*, b_j^*\}$, where

$$(13) \quad p_j^* = t^* [1 + (2j - m - 1)\gamma(m)]$$

$$(14) \quad b_j^* = \hat{e} [1 + N(2j - m)\gamma(m)]$$

$$(15) \quad \gamma(m) = \frac{(1 + a\delta N)\xi}{zN[m + (m - 1)a\delta]}$$

and where

$$(16) \quad \hat{e} = \frac{z}{1 + a\delta N}$$

is the emissions level for the average firm under the Pigouvian tax, t^* .

COROLLARY. The optimal UDBP scheme described in Proposition 1

(a) reduces to the Pigouvian tax when $m = 1$.

(b) is symmetric around the Pigouvian tax; that is,

$$(i) \text{ for } m \text{ odd: } p_{\frac{m+1}{2}} = t^*$$

$$(ii) \text{ for } m \text{ even: } \frac{p_{\frac{m}{2}} + p_{\frac{m}{2}+1}}{2} = t^*$$

This property is illustrated in Figure 2, where t^* passes through the “center” of the UDBP schedule.

(c) is symmetric around the linear function $p^*(e) = \delta e + c$, where

$$c = \frac{z\delta(N - 1)}{1 + a\delta N}$$

That is,

$$\frac{p_j + p_{j+1}}{2} = \delta b_j + c \quad \forall j \in [1, m - 1]$$

and

$$\frac{\tilde{e}_j + \tilde{e}_{m-j+1}}{2} = \hat{e} \quad \forall j \in [1, m] \quad \text{where } \delta \tilde{e}_i + c = p_i$$

This property is also illustrated by the line labeled $p^*(e)$ in Figure 2.

(d) has uniform price steps (above p_1) and uniform price blocks (above b_1); that is,

$$(i) p_{j+1} - p_j = p_{m+1-j} - p_{m-j} \quad \forall j$$

$$(ii) b_{j+1} - b_j = b_j - b_{j-1} \quad \forall j$$

(e) approaches the Pigouvian tax as uncertainty vanishes; that is, $p_j^* \rightarrow t^* \quad \forall j$ as $\omega \rightarrow 0$.

It is important to note that these specific properties of the optimal UDBP scheme are based on the assumption that v_i is distributed uniformly. Departures from this assumption will typically yield different results. To provide a sense of this sensitivity we next present two examples that we have solved numerically, in which v_i is assumed to have a beta distribution with parameters η_1 and η_2 , normalized to have support $[-1,1]$ and zero mean.⁷ In each example the following parameter values are assumed: $a = 0.25$, $\delta = 0.25$, $z = 2$ and $N = 2$. The benchmark for comparison is a uniform distribution with support $[-1,1]$.

First consider a symmetric beta distribution with $\eta_1 = \eta_2 = 3$. The corresponding optimal UDBP scheme (for $m = 3$) is illustrated by the dashed line in Figure 3, alongside the optimal scheme associated with the benchmark uniform distribution (the solid line). Note that the greater central tendency of the beta distribution means that the UDBP scheme under that distribution is flatter and has a narrower central price block than the scheme under a uniform distribution. Both schemes are nonetheless symmetric around the Pigouvian tax (labeled PT in the figure) and around the $p^*(e)$ function identified in the Corollary to Proposition 1.

Next consider a skewed beta distribution with $\eta_1 = 2$ and $\eta_2 = 3$. The median of this distribution is below the mean. The corresponding optimal UDBP scheme (for $m = 3$) is illustrated in Figure 4 (dashed line), alongside the optimal scheme associated with the benchmark uniform distribution (solid line). The key feature of the UDBP scheme in this case is that it is *not* symmetric around the Pigouvian tax nor around the

⁷ The uniform distribution is a special case of the beta distribution where $\eta_1 = \eta_2 = 1$. Higher values of η_1 and η_2 yield a distribution with more central tendency than the uniform distribution.

$p^*(e)$ schedule. The UDBP schedule under the skewed distribution is skewed to the right relative to the symmetric case, reflecting the fact that more than half the firms have an abatement cost below the mean.

These numerical examples point to important limitations of the UDBP scheme: its optimal structure is sensitive to distributional assumptions; and the solution is not straightforward to derive analytically except where v_i has a uniform distribution. These shortcomings can be traced to the sharply discontinuous nature of the pricing scheme associated with the finite number of price blocks. It is therefore natural to ask whether we can do better under a scheme with more continuity. Proposition 2 provides a partial answer.

PROPOSITION 2. If v_i is distributed uniformly with support $[-\xi, \xi]$, then the optimal UDBP scheme has an infinite number of price blocks.

This is an intuitive result, and our conjecture is that it generalizes to *any* distribution – and numerical simulations support that conjecture – but we have not been able to construct a complete proof of the result beyond the case of the uniform distribution. We can however demonstrate the following result.

PROPOSITION 3. The optimal UDBP scheme for any given number of blocks m converges to the continuous price schedule $p^*(e) = \delta e + c$ (as defined in the Corollary to Proposition 1) in the limit as $m \rightarrow \infty$, for *any* distribution of v_i for which the cumulative density is differentiable.

This result tells us that the optimal continuous price schedule is linear. This too is an intuitive result in view of the assumed linearity of MACs and marginal damage in our model. The result – coupled with Proposition 2 and our conjecture regarding its generality – suggest that a continuous linear pricing scheme is superior to a UDBP scheme with a finite number of price blocks. Moreover, such a scheme is easier to derive

than the UDBP scheme under fairly general conditions. We investigate the optimal linear pricing scheme in the next section.

5. THE OPTIMAL ULIMP SCHEME

Recall that our results in section 4 were based on a simplified setting in which $\sigma^2 = 0$ and $\rho^2 = 0$. In this section we relax those assumptions and allow $\sigma^2 > 0$ and $\rho^2 > 0$, together with $\omega^2 > 0$. We do not impose any other distributional assumptions on the random variables.

Suppose all firms face an emissions price schedule of the form

$$(17) \quad p(e_i) = se_i + k$$

where s and k are parameters set by the regulator. That is, the *unit price* on a firm's emissions comprises a constant term k , plus a component that rises with the level of the firm's emissions. In response to this price schedule, firm i chooses emissions

$$(18) \quad e_i(s, k) = \frac{x_i - ak}{1 + as}$$

and aggregate emissions are

$$(19) \quad E(s, k) = \frac{(x - ak)N}{1 + as}$$

The induced values of expected environmental damage and expected aggregate abatement cost are, respectively,

$$(20) \quad E[D(s, k)] = \frac{\delta[N(\omega^2 + N\rho^2) + N^2ak(ak - 2z) + N^2z^2]}{2(1 + as)^2}$$

$$(21) \quad E[C(s, k)] = \frac{Na[s^2(\omega^2 + \rho^2) + k^2 + s^2(z^2 + \sigma^2) + 2ksz]}{2(1 + as)^2}$$

The planning problem is to choose s and k to minimize $E[D(s, k)] + E[C(s, k)]$. The solution to this problem is straightforward to derive, and is described in the following proposition.

PROPOSITION 4. The optimal ULIMP schedule is given by $p(e_i) = s^* e_i + k^*$, where

$$(22) \quad s^* = \frac{\delta(N\rho^2 + \omega^2)}{\omega^2 + \rho^2 + \sigma^2}$$

and

$$(23) \quad k^* = \frac{\delta z[(N-1)\omega^2 + N\sigma^2]}{(1 + N\delta a)(\omega^2 + \rho^2 + \sigma^2)}$$

This optimal ULIMP schedule is illustrated in Figure 5. It has a number of noteworthy properties. First, in the special case where $\sigma^2 = 0$ and $\rho^2 = 0$, we have $s^* = \delta$ and $k^* = c$; the optimal ULIMP schedule coincides with the optimal UDBP scheme as $m \rightarrow \infty$ (recall Proposition 3). Second, in the case with no uncertainty (that is, if $\rho^2 = \omega^2 = 0$), $s^* = 0$ and $k^* = t^*$; the optimal schedule reduces to the Pigouvian tax. Third, if $N = 1$ then $s^* = \delta$ and $k^* = 0$; the optimal schedule coincides with the marginal damage function. Fourth, in the most general case where $N > 1$, $\rho^2 > 0$, $\sigma^2 > 0$ and $\omega^2 > 0$, we have $s^* > 0$ and $k^* < t^*$ (as illustrated in Figure 5), and the following sets of comparative static results obtain:

$$(i) \quad \frac{\partial k^*}{\partial \rho^2} < 0, \frac{\partial s^*}{\partial \rho^2} > 0 \quad \text{and} \quad \frac{\partial k^*}{\partial \sigma^2} > 0, \frac{\partial s^*}{\partial \sigma^2} < 0$$

$$(ii) \quad \frac{\partial k^*}{\partial \omega^2} < (>)0 \quad \text{and} \quad \frac{\partial s^*}{\partial \omega^2} > (<)0 \quad \text{if and only if} \quad \frac{\sigma^2}{N-1} > (<)\rho^2$$

The first set of results capture the essence of the tradeoff underlying the optimal ULIMP schedule: an increasing marginal price schedule has informational advantages but comes with the cost of unequalized MACs when firms are heterogeneous. Accordingly, the optimal price schedule is steeper for higher degrees of aggregate uncertainty (as measured by ρ^2) and flatter for higher degrees of (observable) heterogeneity across firms (as measured by σ^2). The second set of results state that an increase in unobservable heterogeneity across firms (as measured by ω^2) has an ambiguous effect on the optimal price schedule. The reason for that ambiguity is that an increase in ω^2 raises the degree of uncertainty faced by the regulator but also raises the degree of heterogeneity across

firms, thereby creating conflicting forces on the optimal schedule. The net effect depends on the relative size of σ^2 and ρ^2 : if the degree of observable heterogeneity (σ^2) is large, then the additional heterogeneity introduced by a larger value of ω^2 is unimportant relative to the additional uncertainty it introduces, and hence the optimal schedule should be steeper. The converse is true if σ^2 is relatively small.

One additional property of the optimal ULIMP schedule is worth noting here: emissions for firm i under the optimal ULIMP schedule are less than its emissions under the per unit Pigouvian tax if and only if $x_i < z$; otherwise its emissions are greater under the ULIMP schedule. That is, below-average polluters pollute less under the ULIMP scheme than under the Pigouvian tax, while the converse is true for above-average polluters. This property is evident in Figure 5, where the optimal ULIMP schedule, labeled $p^*(e)$, is equal to the Pigouvian tax at the intersection with the average MAC curve, labeled $E[MAC]$.

6. RELATIONSHIP TO OTHER PRICING SCHEMES

An extensive existing literature studies pollution pricing under asymmetric information.⁸ Here we confine consideration to the two papers most closely related to ours: Weitzman (1978) and Roberts and Spence (1976).

6.1 Weitzman (1978)

Weitzman (1978) addresses a very general problem: the design of optimal rewards for economic agents when the benefit function is non-separable in the agents' actions, and

⁸ Formal analysis of environmental regulation under asymmetric information arguably began with Weitzman (1974) where he analyzes the relative merits of price-based regulation and quantity-based regulation under uncertainty about abatement costs. (See Adar and Griffin (1976) for a graphical treatment). The subsequent literature has followed several different directions. One branch has extended the Weitzman (1974) comparative analysis to stock pollutants (for example, Hoel and Karp (2001) and Newell and Pizer (2003)) and to uncertainty about the marginal damage function (for example, Stavins (1996)). A second branch has focused on the design of revelation mechanisms for polluting firms (see for example, Spulber (1988) and Lewis (1996)). A third branch has examined regulatory mechanisms for multiple firms where fees are based on aggregate emissions (for example, Segerson (1988) and McKittrick (1999)). A fourth branch has analyzed dynamic tax adjustment (see Karp and Livernois (1994)). A fifth branch has focused on second-best non-linear pricing schemes and hybrid price-quantity schemes; Roberts and Spence (1976), Weitzman (1978), Yohe (1981) fall into this last category.

both the benefit function and the agents' cost functions are uncertain. He shows that the optimal reward function in this setting comprises a traditional price signal combined with a penalty for deviating from a prescribed quantity target.⁹ In the context of the model in this paper, Weitzman's optimal reward function can be interpreted as a *penalty function*:

$$(24) \quad P_i(e_i) = t^* e_i + \frac{\delta}{2} (\tilde{e}_i - e_i)^2$$

where t^* is the per unit Pigouvian tax (which minimizes *expected* social cost), and \tilde{e}_i is the *expected* first-best emission level for firm i . This hybrid penalty function has the following rationale. If the regulator uses only the Pigouvian tax, based on expected abatement costs, then emissions will be too low relative to first-best if the MAC is lower than expected, and *vice versa*. Conversely, if the regulator uses only a quantity instrument, specifying \tilde{e}_i , then emissions will be too high if the MAC is lower than expected, and *vice versa*.¹⁰ The hybrid scheme in (24) combines elements of both instruments, and is therefore superior to either one alone. It effectively provides a “safety valve” for a quantity target that turns out to be too restrictive *ex post*, and at the same time limits the overshooting that would arise from a tax that turns out to be too lax *ex post*.

Note that the penalty function in (24) corresponds to a marginal price schedule given by

$$(25) \quad p_i^W(e_i) = t^* + \delta(e_i - \tilde{e}_i)$$

where the “W” superscript denotes the Weitzman scheme. Thus, Weitzman's penalty function is equivalent to a LIMP scheme with $s^W = \delta$ and $k_i^W = t^* - \delta\tilde{e}_i$, but it is crucially different from the ULIMP scheme described in section 5 above. In particular, the Weitzman price schedule does not involve uniform treatment across firms; each firm faces a different LIMP schedule, with an intercept tailored to its individual expected abatement cost. This is a crucial distinction because it means that expected MACs are

⁹ Yohe (1981) addresses the same basic problem.

¹⁰ This comparison between prices and quantities is the essence of Weitzman (1974). In that paper he shows that the ranking of the associated expected losses under the two instruments depends on the relative slopes of the MAC and MD functions.

equated across sources under Weitzman's scheme. To see this, note from (18) that emissions from firm i in response to the price schedule in (25) are

$$(26) \quad e_i^w = \frac{x_i - a(t^* - \delta \tilde{e}_i)}{1 + a\delta}$$

Making substitutions for t^* from (8) and for \tilde{e}_i from (10), noting that $\tilde{e}_i = E[e_i^*]$, and calculating the MAC yields

$$(27) \quad MAC_i(e_i^w; \eta, v_i) = t^* + \left[\frac{\delta}{1 + a\delta} \right] (\eta + v_i)$$

Thus, $E[MAC_i(e_i^w; \eta, v_i)] = t^* \quad \forall i$ since $E[\eta] = 0$ and $E[v_i] = 0 \quad \forall i$. This solution is illustrated in Figure 6 for the case of $N = 2$ and $\sigma^2 > 0$. The two firms face different price schedules tailored to the known component of their MAC schedules, and expected MACs are equated. Thus, we face no tradeoff between information benefits and expected abatement costs. (Note however that MACs may not be equated *ex post*, so the Weitzman scheme generally does not achieve first-best).

Now compare the Weitzman scheme with the optimal ULIMP scheme. For simplicity, consider the case where $\rho^2 = 0$. In this case it is straightforward to show that

$$(28) \quad k^* = \left[\frac{\sigma^2}{\omega^2 + \sigma^2} \right] t^* + \left[\frac{\omega^2}{\omega^2 + \sigma^2} \right] \bar{k}^w$$

where \bar{k}^w is the average of the Weitzman intercepts across firms; that is,

$$(29) \quad \bar{k}^w \equiv \frac{1}{N} \sum_{i=1}^N k_i^w = \frac{\sum_{i=1}^N [t - \delta e_i^*]}{N}$$

Thus, k^* is a weighted average of the per-unit Pigouvian tax and the average of the Weitzman intercepts, where the weights reflect the degree of uncertainty ($\omega^2 > 0$) and the degree of observable heterogeneity across firms ($\sigma^2 > 0$) respectively. Since $t^* > \bar{k}^w$, it follows that $k^* > \bar{k}^w$. Note too that $s^* < \delta$. Thus, the optimal ULIMP schedule is flatter and has a higher intercept than the average of the Weitzman schedules.¹¹ This

¹¹ It is not possible to draw such sharp conclusions when $\rho^2 > 0$. The optimal ULIMP is steeper and higher in that case, and cannot be compared unambiguously with the average Weitzman schedule.

reflects the fact that the optimal ULIMP scheme drives a wedge between MACs across firms while the Weitzman scheme does not (in expected terms). The optimal ULIMP scheme must therefore be flatter than the average of the Weitzman schedules. This difference diminishes as ω^2 rises relative to σ^2 because information-related benefits increasingly outweigh the inflation of abatement costs associated with heterogeneity across firms.

The optimal ULIMP schedule and the Weitzman scheme coincide in three special cases. First, if $\omega^2 = 0$ and $\rho^2 = 0$ then both mechanisms reduce to the Pigouvian tax. Second, if $\sigma^2 = 0$ then all firms face the same price schedule under the Weitzman scheme, and that schedule coincides with the optimal ULIMP schedule. The same outcome arises when $N = 1$.

6.2 Roberts And Spence (1976)

Roberts and Spence (1976) – henceforth RS – propose a different sort of hybrid price-quantity scheme, using charges and tradable licenses, where “each [instrument] can protect against the failings of the other. Licenses can be used to guard against extremely high levels of pollution while, simultaneously, effluent charges can provide a residual incentive to clean up more than the licenses required, should costs be low”. (p.194)

The simplest version of the RS mixed scheme works as follows. The regulator issues a total of L licenses. A firm holding l_i licenses pays a penalty equal to $p_2(e_i - l_i)$ if $e_i > l_i$, and receives a subsidy equal to $p_1(l_i - e_i)$ if $e_i < l_i$, where $p_1 \leq p_2$.¹² Thus, in terms of opportunity cost, firm i effectively faces a two-step marginal price function, much like the UDBP scheme illustrated in Figure 2 (but with $m = 2$). However, the RS scheme is importantly different from a simple UDBP scheme because the licenses in their scheme are *tradable* at an endogenous equilibrium price. Thus, the RS scheme functions like a UDBP scheme in which firms are entitled to trade unused portions of the price blocks. Adding this possibility of trade adds an element of administrative complexity to a UDBP pricing scheme, but it also improves its performance. In particular, trade ensures

¹² The original notation used in RS is s and p for our p_1 and p_2 respectively.

that MACs are equated across firms in equilibrium, and this equality holds *ex post*, not just in expectation. Thus, the tradeoff between information benefits and an increase in aggregate abatement cost that arises with a simple UDBP scheme does not arise with the RS scheme.

In an appendix to their paper, RS extend their two-step scheme to one with j types of tradable licenses. Allowing trade in all license types means that the regulated firms effectively act as if they were a single firm facing a marginal price schedule with j steps that approximates the marginal damage schedule. In the limit as $j \rightarrow \infty$, the effective marginal price schedule, in terms of aggregate industry emissions, is a LIMP scheme with $s = \delta$ and $k = 0$.¹³ This limiting scheme achieves first-best results.

It is worth noting that if $N = 1$, then the limit of the generalized RS scheme coincides with the Weitzman (1978) scheme, which in turn coincides with the optimal ULIMP scheme. To see this, recall from (25) that the Weitzman scheme corresponds to a LIMP scheme with $s = \delta$ and $k_i = t^* - \delta \tilde{e}_i$. In the case of a single firm, $t^* = \delta \tilde{e}_i$, so $k_i = 0$. Thus, the Weitzman marginal price function coincides with the marginal damage function, as does the limiting RS price function. Both schemes achieve first-best results in that case.

When $N > 1$ the Weitzman scheme and the RS scheme take different approaches to the information asymmetry problem. The Weitzman scheme fixes individualized quantity targets and allows the marginal price of deviations from those targets to differ across firms (recall Figure 6). In contrast, the RS scheme fixes only an aggregate quantity target and allows firms to trade in quantities; deviations from those individual (endogenous) quantity targets are then priced at the same rate for all firms. This means that MACs are equated across firms *ex post*. In contrast, MACs under the Weitzman scheme are equated only in expectation.

The primary drawback with the RS scheme is that it is generally not possible to derive exact solutions for the optimal subsidy and penalty prices (p_1 and p_2) unless the

¹³ Collinge and Oates (1982) propose essentially the same scheme – a sequence of numbered tradeable permits with differing rental prices – to ensure correct incentives for entry and exit in the polluting

distribution of the *actual* MACs – as opposed to the population distribution from which they are drawn – is known. The problem is the following. Even the simplest form of the RS scheme (with just two steps) has three possible equilibria in the license market, corresponding to where the equilibrium license price is equal to p_1 or p_2 or some value strictly between p_1 and p_2 . The conditional distribution of the MACs in each of these three possible equilibria must be known in order for the optimal values of p_1 , p_2 and L to be calculated, and this in turn generally requires that the sample distribution of the MACs be known. This is a significant information requirement for the regulator.

In the following section we compare the performance of the RS scheme using two steps with the Weitzman scheme and our ULIMP scheme in a simplified setting in which an approximate analytical solution can be found for the optimal RS scheme. However, even in this simplified setting we will see that an exact analytical solution for the optimal RS scheme can be derived only under very special conditions.

7. COMPARATIVE PERFORMANCE

We focus on the special case where $\omega^2 > 0$ but $\sigma^2 = 0$ and $\rho^2 = 0$, and where the v_i 's are drawn from a uniform distribution with support $[-\xi, \xi]$. Recall that this is the same setting examined in section 4, where we derived the optimal UDBP scheme. This setting implies unobservable heterogeneity across firms ($\omega^2 > 0$), but all firms are the same in expectation ($\sigma^2 = 0$), with no industry-wide uncertainty ($\rho^2 = 0$). We have already learned two things about this special case: the optimal UDBP scheme has an infinite number of steps and converges to the ULIMP scheme; and the Weitzman scheme and the ULIMP scheme coincide.¹⁴ Thus, the comparison of interest is between two second-best schemes: the ULIMP/Weitzman scheme and the RS scheme. We begin by deriving the optimal two-step RS scheme (denoted RS2) in this setting.

industry. The outcome is a pricing scheme in terms of aggregate industry emissions that approximates the actual marginal damage schedule.

¹⁴ They may result in unequal MAC's, however, so they do not coincide with the first-best performance of an RS scheme with an infinite number of steps.

7.1 The Optimal RS2 Scheme

The details of the derivation are relegated to the Appendix, but some key elements of the derivation warrant discussion. First, three types of equilibrium can arise under the RS2 scheme, as characterized by the equilibrium price of licenses. In particular, the equilibrium license price can be equal to p_1 , or p_2 , or some value strictly between p_1 and p_2 . The linear functional forms we have specified for MAC and MD mean that the type of equilibrium that arises depends only on the *mean* of the realized v_i 's, denoted \bar{v} , and not on the higher moments of the sample distribution. This simplifies the derivation significantly. Nonetheless, \bar{v} is a random variable from the perspective of the regulator, and its distribution must be known in order for the optimal scheme to be calculated. In the special case where the population distribution of the v_i 's is uniform on $[-\xi, \xi]$, and $N = 2$ (two regulated firms), it can be shown that the distribution of \bar{v} is a symmetric triangular distribution with zero mean and variance equal to $\xi^2 / 6$. In this special case, it is possible to derive an exact analytical solution for the optimal RS2 scheme.

When $N > 2$ the distribution of \bar{v} cannot be found analytically, but the central limit theorem can be invoked to approximate that distribution as a normal distribution. (This of course is true for any population distribution of the v_i 's). This approximation is useful for finding a numerical solution for the optimal scheme, but it is still not possible to find a closed-form analytical solution. However, any normal distribution can be approximated by a symmetric triangular distribution. We use this approach to derive an approximate analytical solution for the RS2 scheme when $N > 2$ (when $N = 2$ our analytical solution is exact). The approximate (or if $N=2$, exact) solution is as follows.

PROPOSITION 5. If $\sigma^2 = 0$ and $\rho^2 = 0$, and v_i is distributed uniformly with support $[-\xi, \xi]$, then the approximately optimal RS2 scheme is given by $\{p_1^{RS}, p_2^{RS}, L^{RS}\}$, where

$$(30) \quad p_1^{RS} = t^* - \frac{\xi \delta \sqrt{2N}}{(3 + a\delta N)}$$

$$(31) \quad p_2^{RS} = t^* + \frac{\xi \delta \sqrt{2N}}{(3 + a\delta N)}$$

$$(32) \quad L^{RS} = \hat{e}N$$

and where

$$(33) \quad \hat{e} = \frac{z}{1 + a\delta N}$$

is the emissions level for the average firm under the Pigouvian tax, t^* . These solutions are exact when $N = 2$.

This optimal RS2 scheme has a number of key properties. First, it is symmetric around the Pigouvian tax (just like the UDBP scheme). Second, it is symmetric around the linear function $p(e) = \delta Ne$. Recall that this is the limiting case of the RS scheme as the number of steps approaches infinity. In contrast, the UDBP scheme is symmetric around the optimal ULIMP. Third, in the case where $N = 2$, the optimal RS2 scheme is steeper than the UDBP schedule in the sense that $p_1^{RS} < p_1^*$ and $p_2^{RS} > p_2^*$. These properties of the RS2 scheme are illustrated in Figure 7.

The steeper profile of the RS2 price schedule relative to the UDBP scheme reflects the fact that the RS scheme does not have to trade off the information benefits of the rising marginal price profile against higher abatement costs, because trade in licenses ensures that MACs are equated in equilibrium. In contrast, the UDBP scheme *does* have to make this tradeoff, and the price profile is flatter as a consequence (so as to reduce the risk of markedly different MACs across firms in equilibrium). It is not possible to prove an unambiguous analytical relationship between the price profiles of the two schemes when $N > 2$, but we suspect that this reflects the approximate nature of our RS2 solution in that case, rather than some breakdown in the logic of the comparison.

We now examine the relative performance of the RS2 and ULIMP schemes. We focus on the ULIMP scheme because we know it is superior to a UDBP scheme with finite steps, both in terms of performance and ease of derivation for the regulator. Our performance comparison has two parts. We first consider an analytical comparison of the minimized expected social cost under the two schemes. We then construct a Monte Carlo simulation to compare their performance in terms of the distribution of actual social cost outcomes.

7.2 Expected Social Cost

We present the comparison between the two schemes in terms of their percentage deviations from the expected first-best social cost. This allows us to conduct comparative statics on the relative performance measure while controlling for the impact of parameter changes on absolute costs. Let SC^{FB} denote the expected value of social cost when emissions are chosen optimally for any realization of the v_i 's. Let SC^{RS2} denote expected social cost under the RS2 scheme, and let SC^{ULIMP} denote expected social cost under the ULIMP scheme. Then we obtain the following results on comparative performance.

PROPOSITION 6. If $\sigma^2 = 0$ and $\rho^2 = 0$, and v_i is distributed uniformly with support $[-\xi, \xi]$, then

$$(a) R^{ULIMP} \equiv \frac{SC^{ULIMP} - SC^{FB}}{SC^{FB}} = \frac{a\delta(N-1)\xi^2}{(a\delta+1)(3Nz^2 + \xi^2)}$$

$$(b) R^{RS2} \equiv \frac{SC^{RS2} - SC^{FB}}{SC^{FB}} = \left(\frac{(a\delta+1)(a\delta+9)N}{(2a\delta+3)^3(N-1)} \right) R^{ULIMP}$$

(c) $R^{ULIMP} - R^{RS2} > 0$ for all $\xi > 0$ and $N \geq 2$; and when $N = 2$, this difference is increasing in a , δ and ξ , and decreasing in z .

The primary message from Proposition 6 is that the RS2 scheme always outperforms the ULIMP/Weitzman scheme. This superior performance of the RS2 scheme stems from the fact that it makes use of increasing marginal prices without sacrificing the equality of MACs in the way that the ULIMP scheme does. It should be noted too that it achieves this superior performance even in its simplest form, with just two steps in the price function. The comparative static results in part (c) indicate that the superiority of the RS scheme is greatest when uncertainty and damage are relatively high (since the ULIMP scheme is steep in that case, which creates greater differences in MACs across firms), and when firms are highly responsive to price differences (as when a is high and/or z is low). While these comparative static results can be proven only when

$N = 2$ (in which case our RS2 solution is exact), it seems reasonable to suppose that they would also hold for larger values of N under the true optimal RS2 scheme.

7.3 Monte Carlo Simulation

In this section we present the results a Monte Carlo simulation in which we compare the realized performance of the two schemes for each of 30,000 draws of the v_i 's from a uniform distribution with support $[-\xi, \xi]$. Our parameter values for the simulation are the same as those used for the numerical examples of the UDBP scheme in section 4: $a = 0.25$, $z = 2$, $\delta = 0.25$, $\xi = 1$, and $N = 2$ (so our solutions for the RS2 scheme are exact). Our qualitative results do not seem to be especially sensitive to changes in these parameters.

Our results are summarized in Figure 8 and Table 1(a). Figure 8 plots relative frequencies of percentage deviations from minimum social cost, for the RS2 scheme, the ULIMP scheme, and the per unit Pigouvian tax (labeled "PT"). These results are broadly consistent with the analytical results on expected cost; the RS2 scheme has a lower frequency of large deviations than either of the other policies. Table 1(a) reports the summary statistics for each policy. Note that while the RS2 scheme performs slightly better on average than the ULIMP scheme, the latter has a somewhat lower variance and a lower maximum deviation. The simple *ex ante* Pigouvian tax is unambiguously the worst performer.

It is noteworthy that the magnitude of the percentage deviations from minimum social cost is generally very small for all three policies. Even the Pigouvian tax has an average deviation of less than 1%, and among 30,000 samples, the maximum deviation is less than 13%. It is tempting to conclude from these numbers that more complicated schemes like the RS2 scheme and even the ULIMP scheme may simply not be worth the administrative trouble relative to the Pigouvian tax. However, the picture looks somewhat different when we examine the components of social cost: abatement cost and damage. Figures 9 and 10 present the relative frequencies of percentage errors in abatement costs and damage (and hence, emissions), respectively. Tables 1(b) and 1(c) report the corresponding summary statistics. A positive deviation in damage means that damage is higher than in the first-best solution; a negative deviation means that damage is lower

than in the first-best solution. Similarly for abatement cost. All three policies have the potential to perform quite badly on these measures, especially with respect to abatement costs. While the assumed policy objective focuses only on total social cost, the political implications of drastically sub-optimal outcomes for emissions and abatement costs are likely to be an issue of concern to regulators. The RS2 scheme still outperforms the ULIMP scheme on these measures in terms of averages, but it finishes second behind the ULIMP in terms of variability. The Pigouvian tax finishes last on both fronts.

On balance, our results suggest that the ULIMP scheme does not perform significantly worse on average than the RS2 scheme, and produces less variability in performance outcomes. The optimal ULIMP scheme can also be derived analytically under much weaker distributional assumptions than are required for an exact analytical solution for the RS2 scheme. Coupled with the fact that the RS2 scheme requires an additional administrative layer associated with license trading, these considerations may make the ULIMP scheme a more appealing choice to regulators tasked with implementing a simple but reasonably effective policy.

8. CONCLUSION

Most policies for pricing pollution under asymmetric information proposed in the literature to date are rarely – if ever – used in practice. This is likely due to their complexity. We have investigated the scope for using somewhat simpler policies that are more closely related to pricing schemes already used by regulators in many jurisdictions. These schemes involve a schedule of increasing marginal prices applied uniformly across firms. Under fairly restrictive assumptions, we have derived the optimal form of the uniform discrete block pricing (UDBP) scheme. We have also derived the optimal limiting case of the UDBP schedule (with an infinite number of price blocks) as a linear increasing marginal price schedule with uniform treatment across firms (ULIMP). The latter can be derived directly under less restrictive assumptions and is arguably no more complicated to implement in practice than the UDBP with a finite number of steps.

The optimal ULIMP scheme strikes a balance between the information-related benefits of increasing marginal prices on one hand, and an increase in aggregate abatement cost on the other, due to the non-equalization of MACs across firms. In

particular, the optimal schedule is steeper for larger aggregate uncertainty about MACs, and flatter for more observable heterogeneity across firms.

We have also compared our pricing schemes with the more sophisticated schemes proposed by Weitzman (1978) and Roberts and Spence (1976), which can achieve information-related benefits without sacrificing cost-effectiveness (at least in expected terms). While clearly superior in theoretical terms, these schemes are more complex from a practical perspective, requiring different firms in the same regulated industry to face different price schedules (as in Weitzman) or a mix of discrete block pricing and license trading (as in Roberts and Spence). Moreover, the informational requirements for the optimal Roberts and Spence scheme – with respect to the distribution of the actual MACs – are significant. Our comparative performance results confirm the superiority of the Roberts and Spence scheme in terms of expected cost, even in its simplest two-step form, but the ULIMP scheme does not perform significantly worse, and produces less variable results.

While the relative simplicity and familiarity of the price schemes we have proposed here may have some appeal to regulators, it must be noted that even these schemes are complicated to design in practice. The regulator needs information about the mean and variance of the unknown parameters, and must make some conjecture about the shape of MACs. The linear functional forms we have assumed here may be a reasonable approximation to reality in many circumstances, but not in others. In such cases the policy design problem is considerably more complicated, and any actual pricing scheme is likely to be only a rough approximation to an optimal scheme, derived not from pure analytics but from numerical simulations. We view the analytical results we have derived here as a potentially useful point of departure.

APPENDIX

Note on Proposition 1.

We solved equations (A2) to (A6) below for $m=2$, $m=3$, $m=4$ and $m=5$ (using *Maple Version 8*) and from those solutions derived by induction the general solution reported in Proposition 1. The *Maple* code is available from the authors.

Proof of Proposition 2.

It is straightforward to show that total social cost under the optimal UDBP scheme is given by

$$(A1) \quad SC(m)_{UDBP} = \frac{\delta N[3Nz^2(1+a\delta) + \xi^2(1+a\delta N)]}{6(1+a\delta)(1+a\delta N)} + \frac{a\delta^2 N^2 \xi^2}{6(1+a\delta)[m(1+a\delta) - a\delta]^2}$$

This is decreasing in m . Thus, total social cost is minimized by setting m as large as possible.¹⁵

Proof of Proposition 3.

It is instructive to first prove this result in the context of the special case where $f(v)$ is a uniform distribution. Recall that equations (13) – (16) describe the optimal UDBP scheme in that case. In the limit as $m \rightarrow \infty$, a different price is set for every level of emissions. So from (14) set $b_j^* = e$ and solve for $j(e)$. Then substitute $j(e)$ for j in p_j^* from (13) to yield $p^*(e)$.

Now consider the general proof. Let $f(v)$ denote the distribution of v on support $[v_L, v_H]$ and let $F(v)$ denote the associated cumulative density. We assume that $F(v)$ is twice continuously differentiable. Then the first-order conditions for a minimum are given by equations (A2) to (A5):

$$(A2) \quad [F(b_1 + ap_1 - z) - F(v_L)](p_1 + \delta z - a\delta(N-1)E[e]) + \int_{v_L}^{b_1 + ap_1 - z} v f(v) dv = 0$$

$$(A3) \quad [F(v_H) - F(b_{m-1} + ap_m - z)](p_m + \delta z - a\delta(N-1)E[e]) + \int_{b_{m-1} + ap_m - z}^{v_H} v f(v) dv = 0$$

For $j = 2$ to $m - 1$:

$$(A4) \quad [F(b_j + ap_j - z) - F(b_{j-1} + ap_j - z)](p_j + \delta z - a\delta(N-1)E[e]) \\ + \int_{b_{j-1} + ap_j - z}^{b_j + ap_j - z} vf(v)dv = 0$$

For $j = 1$ to $m - 1$:

$$(A5) \quad a\delta[F(b_j + ap_{j+1} - z) - F(b_j + ap_j - z)](b_j + z + (1-a)(N-1)E[e]) \\ + \int_{b_j + ap_j - z}^{b_j + ap_{j+1} - z} vf(v)dv = 0$$

where

$$(A6) \quad E[e] = \int_{v_L}^{b_1 + ap_1 - z} (z + v - ap_1)f(v)dv \\ + \sum_{j=2}^{m-1} \left[\int_{b_{j-1} + ap_{j-1} - z}^{b_{j-1} + ap_j - z} b_{j-1}f(v)dv + \int_{b_{j-1} + ap_j - z}^{b_j + ap_j - z} (z + v - ap_j)f(v)dv + \int_{b_j + ap_j - z}^{b_j + ap_{j+1} - z} b_jf(v)dv \right] \\ + \int_{b_{m-1} + ap_m - z}^{v_H} (z + v - ap_m)f(v)dv$$

In equation (A4), divide throughout by $(b_j - b_{j-1})$ and take the limit as $b_{j-1} \rightarrow b_j$ to obtain

$$(A7) \quad F'(b_j + ap_j - z)(p_j + \delta z - a\delta(N-1)E[e]) + \lim_{b_{j-1} \rightarrow b_j} \int_{b_{j-1} + ap_j - z}^{b_j + ap_j - z} vf(v)dv = 0$$

Similarly, in equation (A5), divide throughout by $a(p_j - p_{j+1})$ and take the limit as $p_{j+1} \rightarrow p_j$ to obtain

$$(A8) \quad F'(b_j + ap_j - z)(\delta b_j + \delta z + \delta(1-a)(N-1)E[e]) + \lim_{p_{j+1} \rightarrow p_j} \int_{b_j + ap_j - z}^{b_j + ap_{j+1} - z} vf(v)dv = 0$$

It then follows from (A7) and (A8) that

¹⁵ It is worth noting that setting $m=1$ in (A1) yields the total SC under the Pigouvian tax.

$$(A9) \quad p_j + \delta z - a\delta(N-1)E[e] = \delta b_j + \delta z + \delta(1-a)(N-1)E[e]$$

which reduces to

$$(A10) \quad p_j = \delta b_j + \delta(N-1)E[e]$$

Now in (A6) take the limits as $b_{j-1} \rightarrow b_j$ and $p_{j+1} \rightarrow p_j$ and substitute for p_j from (A10) to obtain

$$(A11) \quad E[e] = \int_{v_L}^{v_H} (z + v - a\delta b_j - a\delta(N-1)E[e])f(v)dv$$

Now set $b_j = e$, note that $E[v] = 0$, and then solve (A11) to obtain

$$(A12) \quad E[e] = \frac{z}{1 + a\delta N}$$

Then substituting (A12) into (A10) we obtain

$$(A13) \quad p(e) = \delta e + \frac{z\delta(N-1)}{1 + a\delta N}.$$

Sketch Proof of Proposition 5.

Let q denote the equilibrium price of licenses. Without loss of generality, suppose the regulator issues an equal number of licenses, $l = L/N$, to each firm. There are three equilibrium types: (1) $q = p_1$; (2) $q = p_2$; and (3) $p_1 < q < p_2$. First consider

equilibrium type (3). Faced with license price $q \in (p_1, p_2)$, firm i emits

$e_i(q) = z + v_i - aq$. Firms for whom $e_i(q) < l$ are license suppliers, and sell

$s_i(q) = l - (z + v_i - aq)$ licenses. Firms for whom $e_i(q) > l$ are license buyers, and

demand $d_i(q) = z + v_i - aq - l$. Equilibrium occurs where

$$(A14) \quad \sum_{i=1}^K s_i(q) = \sum_{i=K+1}^N d_i(q)$$

where K is the number of seller firms. Solution of (A14) yields

$$(A15) \quad \tilde{q} = \frac{z + \bar{v} - l}{a}$$

where $\bar{v} = \sum_{i=1}^N v_i / N$. The associated level of aggregate emissions is L . This outcome

can be an equilibrium only if $p_1 < \tilde{q} < p_2$. If $\tilde{q} < p_1$ then no firm will be willing to sell

licenses at $q = \tilde{q}$ since it can instead receive a per unit payment of $p_1 > \tilde{q}$ from the regulator for reducing emissions below its license holdings. In that case we have equilibrium type (1), where $q = p_1$. The equilibrium level of emissions in that case is

$$(A16) \quad E_1 = \sum_{i=1}^N e_i(p_1) = N(z + \bar{v} - ap_1) < L$$

Conversely, if $\tilde{q} > p_2$ then no firm will be willing to buy licenses at $q = \tilde{q}$ since it can instead pay the regulator $p_2 < \tilde{q}$ per unit to emit beyond its license holdings. In that case we have equilibrium type (2), where $q = p_2$. The equilibrium level of emissions in that case is

$$(A17) \quad E_2 = \sum_{i=1}^N e_i(p_2) = N(z + \bar{v} - ap_2) > L$$

We can now characterize the complete equilibrium in terms of a partition of the interval $\bar{v} \in [-\xi, \xi]$, since the sample mean \bar{v} can lie anywhere within this interval. In particular, we have equilibrium type (1) if and only if $\bar{v} < l + ap_1 - z$; equilibrium type (2) if and only if $\bar{v} > l + ap_2 - z$; and equilibrium type (3) if and only if $l + ap_1 - z < \bar{v} < l + ap_2 - z$.

To make further progress it is necessary to specify the distribution of \bar{v} . If the population distribution of the v_i 's is uniform on $[-\xi, \xi]$, then it can be shown that the mean (\bar{v}) of a sample of size $N = 2$ has a symmetric triangular distribution with support $[-\xi, \xi]$, zero mean and a variance equal to $\xi^2 / 6$. If $N > 2$ then we must invoke the central limit theorem. In particular, the distribution of \bar{v} is approximately normal, with zero mean and a variance equal to $\xi^2 / 3N$. This distribution can in turn be approximated by a symmetric triangular distribution with support $[-\theta, \theta]$, zero mean and a variance equal to $\theta^2 / 6$, where $\theta = \xi\sqrt{2}/\sqrt{N}$. It is then possible to construct a conditional density for v_i for each equilibrium type, conditional on \bar{v} satisfying the restriction corresponding to that particular equilibrium type. In each case that conditional density is a truncated triangular distribution whose distribution can be derived exactly. It is then straightforward to calculate expected social cost as a function of $\{p_1, p_2, L\}$. Choosing

these policy parameters to minimize expected social cost yields the solution reported in Proposition 5.

REFERENCES

- Adar, Z. and J.M. Griffin (1976), "Uncertainty and the choice of pollution control instruments", *Journal of Environmental Economics and Management*, 3, 178-188.
- Blair, B.F. (1995), "Sharp regulatory performance with blunt instruments", *Economics Letters*, 48, 201-206.
- Collinge, R.A. and W.E. Oates (1982), "Efficiency in pollution control in the short and long runs: a system of rental emission permits", *Canadian Journal of Economics*, 15, 346-354.
- Dasgupta, P., P. Hammond and E. Maskin (1980), "On imperfect information and optimal pollution control", *Review of Economic Studies*, 47, 857-860.
- Hoel, M. and L. Karp (2001), "Taxes and quotas for a stock pollutant with multiplicative uncertainty", *Journal of Public Economics*, 82, 91-114.
- Kaplow, L. and S. Shavell (2002), "On the superiority of corrective taxes to quantity regulation", *American Law and Economics Review*, 4, 1-17.
- Karp, L. and J. Livernois (1994), "Using automatic tax changes to control pollution", *Journal of Environmental Economics and Management*, 27, 38-49.
- Lewis, T.R. (1996), "Protecting the environment when costs and benefits are privately known", *Rand Journal of Economics*, 27, 819-847.
- McKittrick, R.R. (1999), "A Cournot mechanism for pollution control under asymmetric information", *Environmental and Resource Economics*, 353-363.
- Newell, R.G. and W.A. Pizer (2003), "Regulating stock externalities under uncertainty", *Journal of Environmental Economics and Management*, 45, 416-432.
- Roberts, M.J. and M. Spence (1976), "Effluent charges and licenses under uncertainty", *Journal of Public Economics*, 5, 193-208.
- Segerson, K. (1988), "Uncertainty and incentives for nonpoint pollution control", *Journal of Environmental Economics and Management*, 15, 87-98.
- Stavins, R.N. (1996), "Correlated uncertainty and policy instrument choice", *Journal of Environmental Economics and Management*, 30, 218-232.
- Spulber, D.F. (1988), "Optimal environmental regulation under asymmetric information", *Journal of Public Economics*, 35, 163-181.

- Watson, W.D. and R.G. Ridker (1984), "Losses from effluent taxes and quotas under uncertainty", *Journal of Environmental Economics and Management*, 11, 310-326.
- Weitzman, M.L. (1974), "Prices vs. quantities", *Review of Economic Studies*, 41, 477-491.
- Weitzman, M.L. (1978), "Optimal rewards for economic regulation", *American Economic Review*, 68, 683-691.
- Yohe, G.W. (1981), "Should sliding controls be the next generation of pollution controls?", *Journal of Public Economics*, 15, 251-267.

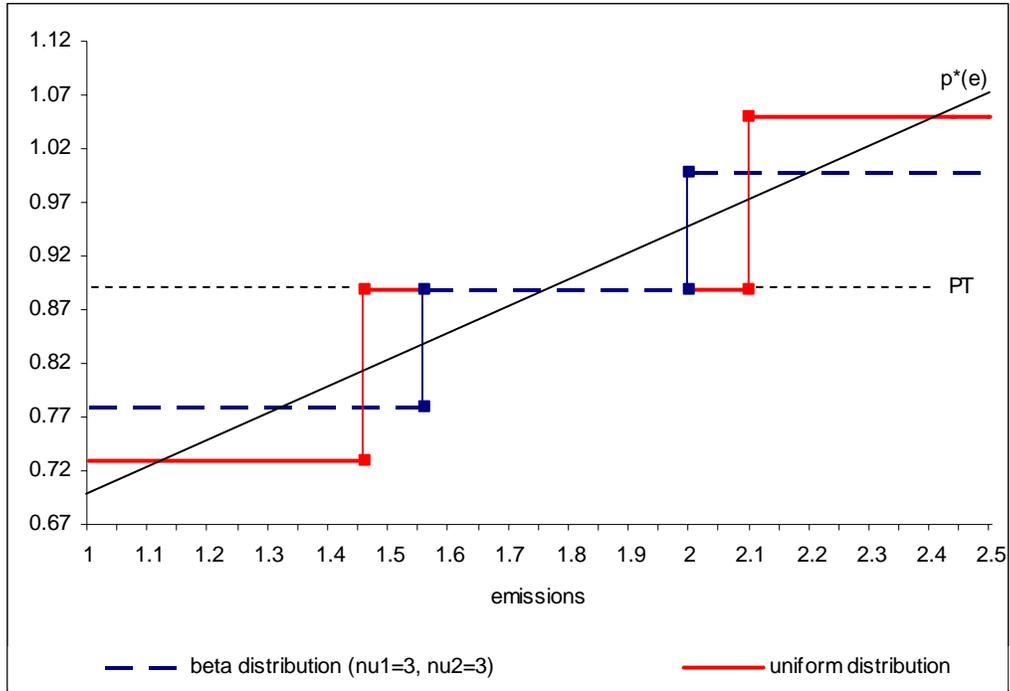


FIGURE 3

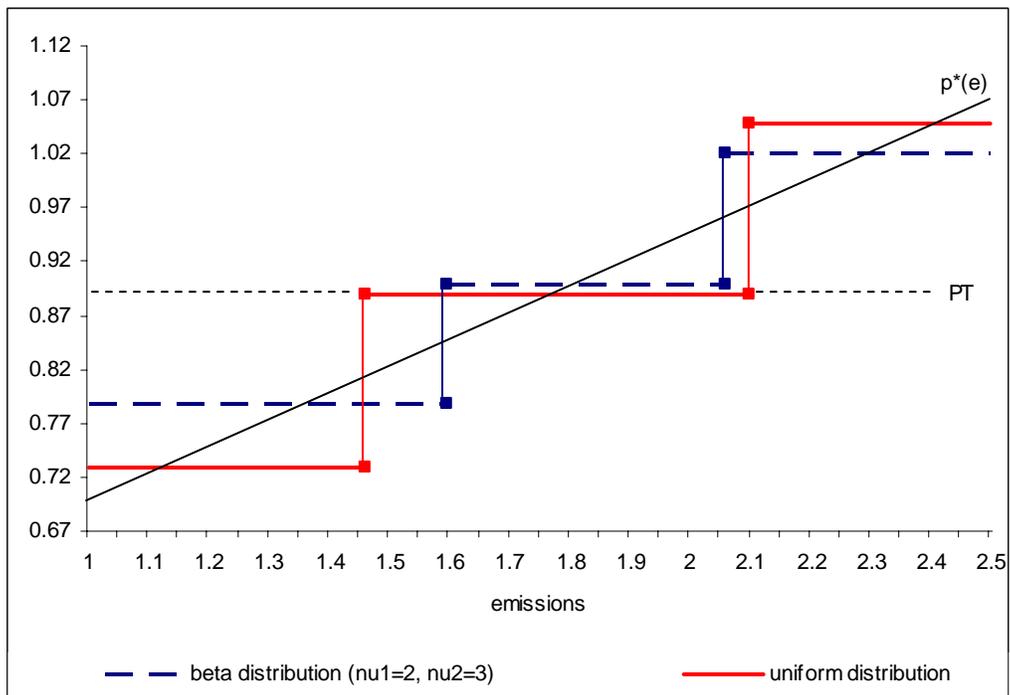


FIGURE 4

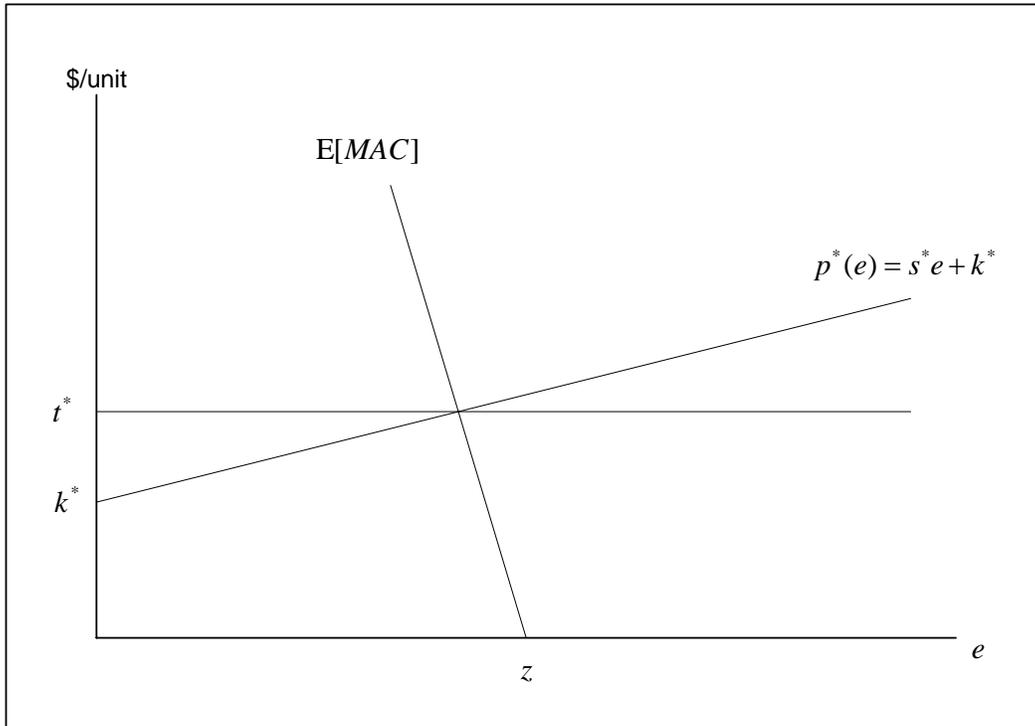


FIGURE 5

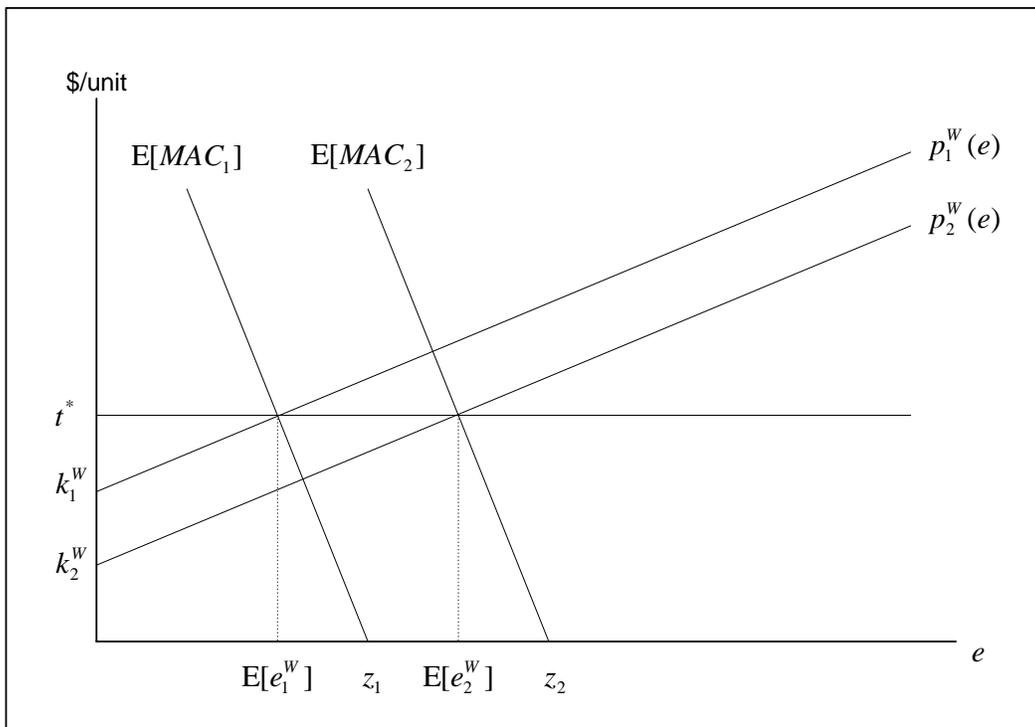


FIGURE 6

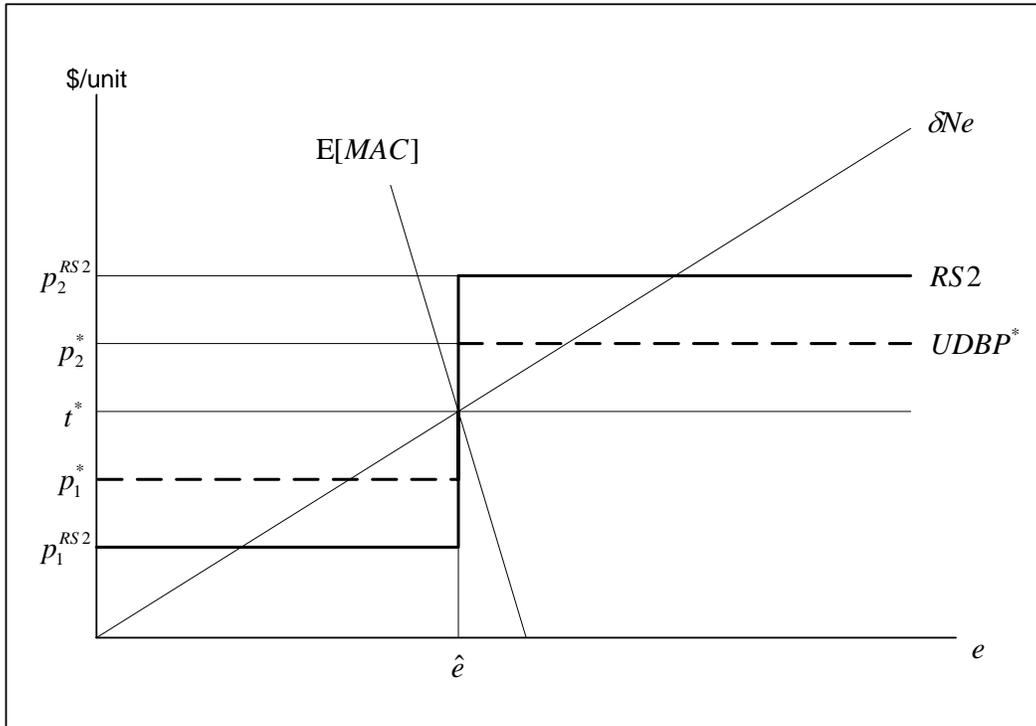


FIGURE 7

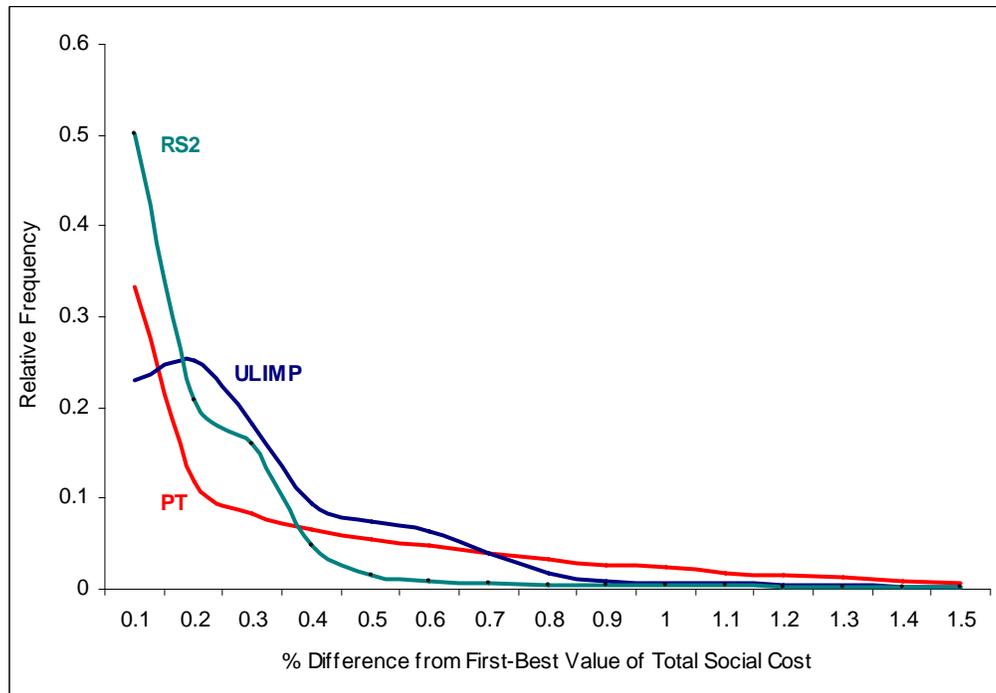


FIGURE 8

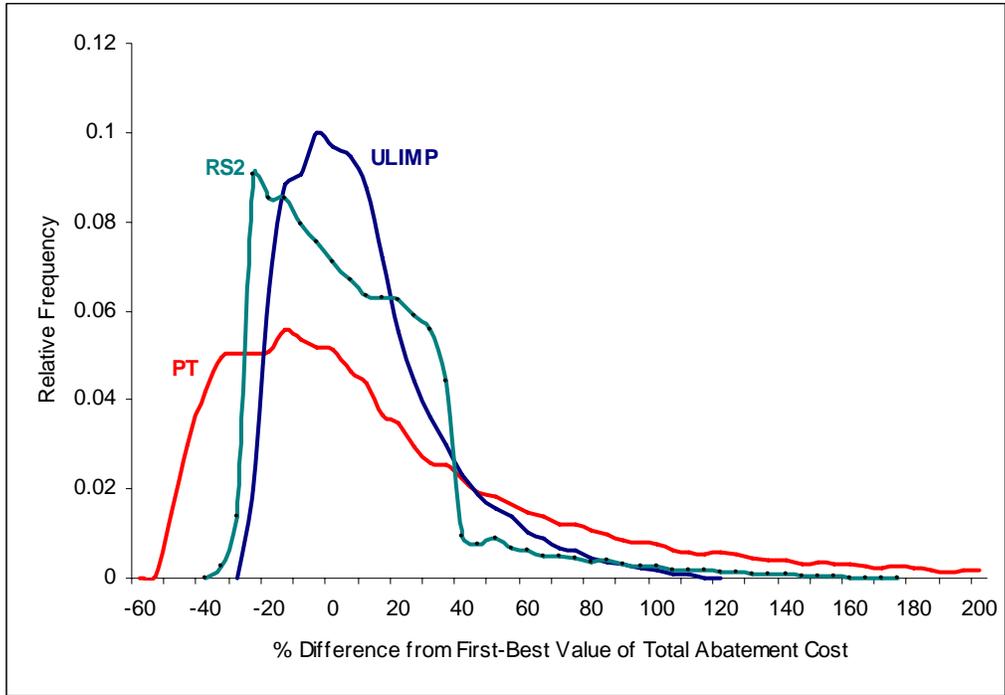


FIGURE 9

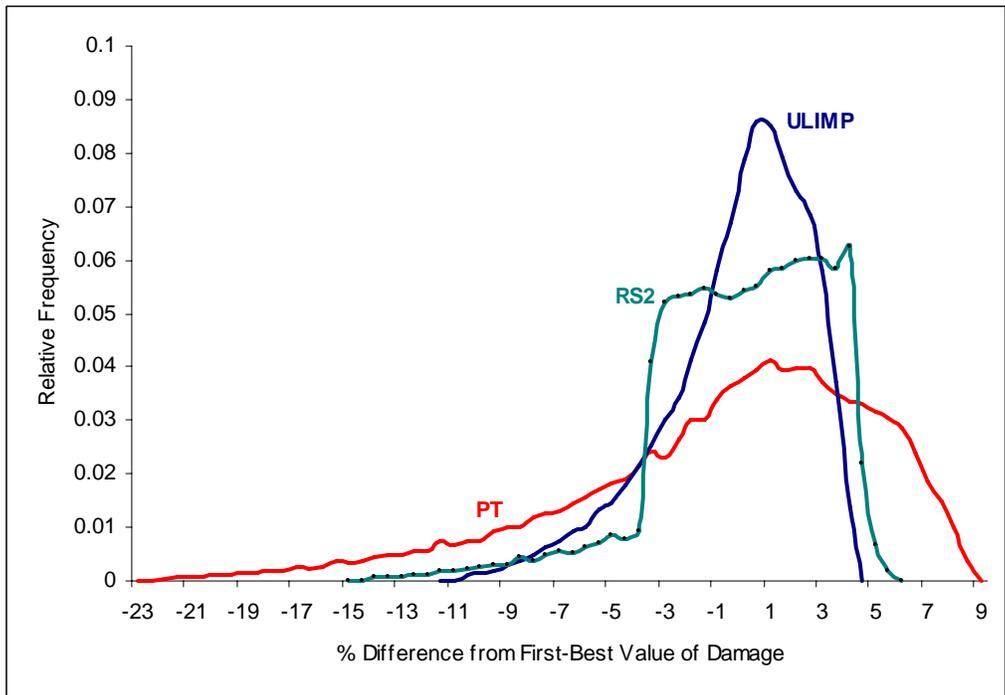


FIGURE 10

(a) % Difference from First-Best Value of Total Social Cost

	PT	ULIMP	RS2
maximum	12.2388	2.7103	4.9831
minimum	0	0	0
mean	0.7137	0.2955	0.2065
std deviation	1.3322	0.3005	0.4185

(b) % Difference from First-Best Value of Total Abatement Cost

	PT	ULIMP	RS2
maximum	295.8095	114.8117	166.1423
minimum	-55.4686	-28.8602	-37.9945
mean	14.7937	6.7765	4.3125
std deviation	55.0174	23.8841	27.7239

(c) % Difference from First-Best Value of Damage

	PT	ULIMP	RS2
maximum	8.4900	3.9522	5.3847
minimum	-23.2076	-11.3023	-15.1618
mean	-1.0463	-0.5146	-0.3068
std deviation	5.7476	2.7339	3.1645

TABLE 1