Three Essays on the Theory and Practice of Environmental Standard-Setting.

Anthony G. Heyes
Department of Economics
McGill University, Montréal
July, 1993.

A Thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment of the requirements of the degree of Doctor of Philosophy.

ISBN 0-315-91722-9

Shortened Title:

"Three Essays on Environmental Standard-Setting"

(Anthony G. Heyes)

Abrégé

Cette thèse présente trois études portant sur las réglementation des externalités. Dans le premier chapitre, nous utilisons un modèle à deux étapes afin d' analyser l'application des réformes politiques environnementales lorsque la technologies d'inspection est imparfaite. Nous démontrons que lors de l'évaluation de ces réformes, il est important de differencier entre deux types de violateurs. Ces violateurs se distinguent par leurs réactions vis-à-vis la reglementation. Dans le deuxieme chapitre, nous démontrons l'inefficacité qu'engendre l'utilisation du critère ALARA lorsque l'offre de nouvelles technologies protégant l'environnement est endogène. L'on charactérise la différence entre ALARA et le critère optimal. On démontre que cette difference dépendra de si la source de changement technologique provient de la R.&D. ou de l'expérience de la firme, ou encore, d'une combinaison de ces deux sources d'apprentissage. Dans le troisième chapitre, nous examinons les externalités qui cont créées par un type particulier de règlementation. Nous démontrons l'existence d'une classe de gains qu'engendrerait l'auto-réglementation de cette industrie. Chacun de ces chapitres discute le cas d'une industrie. Chacun de ces chapitres discute le cas d'une industrie spécifique bien que les résultats visent une application plus générale.

Abstract

This thesis presents three studies on the regulation of externalities. In the first chapter we use a two-period model to analyse the reform of environmental policy when enforcement is incomplete. We show that in evaluating such reforms it is important to differentiate between two types of non-compliers which react to regulatory reform in qualitatively distinct ways. In Chapter 2 we show the problems caused by using the 'ALARA' criterion when the supply of new technologies of environmental protection is endogenous. The difference between the ALARA rule and the socially optimal rule is characterised and it is shown that the difference depends upon whether the source of technical change is R&D effort or learning-by-doing. In the third chapter we examine the externalities which are created by a particular type of regulation, and show the gains from self-regulation which are thus generated. Though each chapter discusses a particular industry or class of industries, the results obtained are argued to have more general validity.

Table of Contents.

(i)-(vii) Introduction.

Chapter 1: Environmental Uncertainty and Environmental Enforcement: The Case of the U.S.E.P.A.

1 (0) Introduction.
5 (1) Selective Survey: Models of Non-Compliant Behaviour.
11 (2) A Two-Stage Characterisation of the Monitoring Process.
34 (3) Two-Stage Enforcement, Tightening Referral Practices and Non-Compliant Behaviour.
45 (4) Industry Attitudes Towards 'Inspectability': The Potential Endogeneity of Inspection Noise.
58 (5) Conclusions.
60 Appendix to Chapter 1.

Chapter 2: Environmental Standard-Setting With Endogeneous Technical Change: A Dynamic Model and Application to the Nuclear Electricity Sector.

61 (0) Introduction.
64 (1) The 'ALARA' Rule in National and Supra-National Regulation of the Nuclear Industry.
68 (2) Environmental Standard-Setting and Innovation: A Selective Survey.
71 (3) A Static Model of Safety Regulation.
76 (4) A Simple Two-Period Model of Safety Regulation With R&D.
95 (5) A Dynamic Model of Safety With Learning-by-Doing: On Learning to Operate a Nuclear Reactor Safely.
114 (6) A Model With R&D and Learning: The Case of 'Experience Directed Research and Development' (EDR&D).
132 (7) Conclusions: Reactive Regulation and the Sources of Technical Knowledge.
135 Appendix to Chapter 2.

Chapter 3: Market Structure, Regulatory Externalities and the Over-loading of Trucks: An Industry Perspective.

141 (1) Introduction to the Problem.
145 (2) A Stylised Model of the Truck-Loading Decision.
149 (3) Market Share and Loading.
160 (4) On the Benefits From Industry Self-Regulation.
170 (5) Conclusions.

172 Conclusions
173 Bibliographies

introduction.

In this thesis we explore three aspects of the regulation of externalities. An externality is created when some fraction of the costs or benefits generated by the activities of one agent are incident upon some other agent or agents. It is well known that, by driving a wedge between social and private evaluations of marginal costs and/or benefits, they serve to distort individual incentives and constitute a significant form of market failure. It is straightforward to show that the presence of externalities reduces social welfare and provides a rationale for policy intervention.

The study of externalities is particularly prevalent in the field of environmental economics - reflecting the frequency with which they are the basis for policy action on environmental issues. The familiar story of the smoke-belching factory and the neighbouring laundry captures, in a stylised way, the archetypel environmental problem. One agent, by his actions, damages the environment of another whose utility is thereby reduced.

In Chapters 1 to 3 we examine three variations on the factory/laundry scenario. In each case simple economic theory provides a framework for the analysis of topical environmental policy issues.

In Chapter 1 we consider the problem faced by a regulatory agency of enforcing compliance with an environmental standard when there is evidentiary uncertainty. We assume that the agency is able to conduct two types of checks: 'inspections' and 'audits'. The former are non-exhaustive and yield only preliminary evidence regarding the compliance status of the regulated firm, whereas the latter are exhaustive. The budget-constrained agency uses evidence from randomly conducted inspections in order to target which firms are to be audited. An inspection generates a noisy (but unbiased) estimator of the true status of the firm and the agency sets some 'audit trigger' such that an inspection estimate in excess of that trigger causes the inspected firm to be audited.

The interest in the study of a 'two-stage' enforcement régime is motivated by the practices of the Environmental Protection Agency in the US, and a number of other environmental agencies. In discussion of enforcement issues it is usual to draw a distinction between compliers and violators. The divergence of the audit trigger from the underlying standard implies that we can distinguish between to two qualitatively distinct violator 'types': 'serious violators' and 'non-serious violators'. Firms in both groups are in violation of the underlying environmental standard, but only those in the former are so far in violation that an inspection would, in most cases, trigger an audit.

The distinction <u>between</u> violators is shown to be important because serious and non-serious violators will (as is shown) react in different ways to various kinds of regulatory reform. The efficacy of these reforms can, then, be expected to depend upon the 'anatomy' of the regulated population. The nature and implications of this dependence are examined at length.

The two-stage enforcement framework is also used to rationalise the coexistence over the last decade of (i) a tightened audit policy by the EPA in the enforcement of air pollution standards with (ii) a growing 'polarisation' of violators. This is compatible with our model. We show that a tightening of audit referral practices induces an improvement in compliance from non-serious violators, but the opposite from those firms initially in serious violation: the bad get better while the very bad get worse.

In addition, we use the two-stage enforcement construct to reassess the conventional wisdom that a non-compliant firm will always benefit from an increase in the 'noisiness' of the inspection technology used by the enforcement agency. We demonstrate that such an assertion is invalid within the framework we employ. In general, the serious violator will favour an increase in inspection uncertainty, whereas the non-serious violator will prefer the reverse (though both of these signs can, in some circumstances, be reversed). The implications of this for the incentive facing different types of firm to invest in 'inspection-proofing' of their operations or, conversely, to help the inspection teams in their work is noted.

In Chapter 2 we consider the problems associated with the determination of optimal environmental standards in an explicitly dynamic framework in which the development of new technologies for the protection of the environment is endogeneous. The principle that environmental risk be kept 'As Low As Reasonably Achievable' (ALARA) is deep-rooted in the regulatory principles and practices of a number of national and supra-national regulatory bodies. We focus on the application of the ALARA-rule in the setting of safety requirements for nuclear power reactors, though our results have much broader relevance.

In the context of nuclear reactor regulation, the ALARA-rule dictates that at any moment the probability of an accident at a nuclear plant should be set at that level at which the social marginal costs and benefits of additional increments of risk-reduction are exactly balanced. ALARA is a particular case of technology-based standards - the stringency of safety regulations at any particular time is determined by the current state of technology. Whilst this flexibility is a supposed strength, it may also be its critical weakness. The 'moving goal-post' regulations which it generates in a dynamic environment will, we argue, inhibit the supply of new techniques of externality control. The incentive facing the regulated industry to develop safety equipment is blunted by the anticipation of the upward ratchetting of requirements which the ALARA criterion inevitably dictates must follow such a development. These inter-temporal inefficiencies may swamp the intra-temporal efficiency gains which the flexibility yields, particularly in a sector such as the nuclear power industry in which the scope for rapid technological advance is so great.

In Chapter 2 we develop a simple two-period model which is used to examine the regulated industry's incentives to devise and commercialise new safety technologies under different regulatory régimes. The dynamic behaviour of the industry is characterised. Optimal standard-setting subject to this dynamic incentive compatibility constraint is examined in some detail, and the way in which regulators should depart from the ALARA rule characterised. We consider the cases in which the source of technological improvement is, in turn: (i) R&D effort, (ii) Learning by Doing and (iii) Experience-Directed R&D ('EDR&D'). The latter is a hybrid of the other two, in which R&D is the source of technological change, but productive experience serves to 'target' that R&D. We argue that EDR&D may be particularly applicable to the context of nuclear safety technologies. We find that the nature of the optimal departure from ALARA depends qualitatively upon the source of the technical progress and upon the extent to which the regulatory agency is able to commit to future standards. This dependence is characterised and the familiar analysis of 'technology-forcing' is reassessed and its inadequacies highlighted.

In Chapter 3 we consider the problem of the overweighting of trucks by truck operators. The overweighting of trucks can do considerable damage to the man-made environment. By causing degradation to the pavement surface, overladen trucks impose costs on other road-users in terms of increased travel times, and decreased comfort and road safety. By increasing the amount of money which must be spent on road maintenance the behaviour also imposes a burden on the public purse and thence upon society in general. Though the action of overweigthing causes, in this way, a 'real' externality, it is 'regulatory' externalities which are the focus of our analysis in Chapter 3.

We construct a simple model of the enforcement of a truck-weight limit in which each trucking firm chooses how often to overload trucks in its fleet. We assume (and then motivate, in a number of ways) that the intensity of the regulatory agency's verification programme in any geographical segment of the market increases when the rate of violation amongst the truck population in that segment increases. In this case the firms in the industry inflict negative regulatory externalities on each other - by overloading one of its trucks Firm A inflicts a cost on Firm B in the form of a tougher regulatory environment. In this context we demonstrate that a more concentrated industry will violate less often than a less concentrated one and that a small firm will violate more often than a large one. We discuss at some length the possibility of industry self-regulation, whereby industry members agree to codes of practice on weight limits. We demonstrate that there will, in general, be gains to all members of the industry from the succesful implementation of a self-regulatory agreement to reduce the frequency with which weight limits are violated. Such an agreement would serve to internalise a class of regulatory externalities, to the benefit of all. We outline problems with the implementability and sustainability of cooperative, self-regulatory action by the industry and discuss the mechanisms which could be used to improve its chances of success.

To avoid confusion, the reader should bear in mind that each of the three chapters is essentially self-contained. Though they share a common theme, namely the regulation of externalities, the analysis presented in each should be viewed as independent of the analysis in each other. In particular, notation and assumptions made in one chapter are <u>not</u> carried over to the next. The list of references associated with each chapter is presented separately at the end of that chapter. Selective surveys of relevant, previous work in each field are presented as the first sections of each chapter.

Evidential Uncertainty and Environmental Enforcement: The Case of the U.S.E.P.A.

(0) Introduction.

Mainstream discussion of optimal standard-setting (in various policy contexts) implicitly assumes that the standards set will be fully complied with. There is, however, a growing literature which seeks to examine how the existence of partial compliance will affect the standard-setting problem. It is to this debate that we hope to add some novel results. More importantly, we examine the possible environmental and budgetary impacts of different types of reform in United States environmental legislation given the enforcement régime operated by the regulatory agency, the Environmental Protection Agency (U.S.E.P.A., or EPA forthwith). The EPA is the principal institution with responsibility for executing US environmental legislation. In particular, it has responsibility for enforcement of the US Clean Air Act (1970) and the US Federal Water Pollution Control Act (1972). Our results suggest that the relationships between the intentions of reformers and the consequences of the suggested reforms may be less transparent than is popularly supposed. We examine the factors which combine to determine these relationships.

Analysis of changes in the behaviour of *non-compliant* firms is of policy significance. The extent to which the real world is in fact populated by non-compliant agents will dictate the importance, in any given context, of explicitly accounting for the responses of non-compliers to regulatory change.

In the context of U.S. environmental regulations there is evidence of widespread non-compliance; in some cases compliance may even be the exception rather than the rule. In the case of air pollution, for example, the appropriate regulations are laid out in the US Clean Air Act (and associated revision documents) and enforced by the EPA. Magat (1990) and Wasserman (1984) estimate that compliance rates with respect to certain classes of pollutants may be as low as 20 %. The frequency and severity of non-compliance is substantial. Actual compliance rates may even be lower than suggested by estimates from surveys of this type, which are based on the data-bases of the regulatory agency. Sophisticated polluters can increase their emissions when inspectors are unlikely to visit: a 1979 follow-up study by the U.S. General Accounting Office (U.S.G.A.O. (1979)) indicated that about 25% of inspected sources which had been deemed 'in compliance' were in fact significantly exceeding their overall emissions limits.

In this paper we construct a simple two-stage model of enforcement. A first round of random inspections generates uncertain "evidence" regarding the compliance status of those firms inspected. The "evidence" is then used to target a second round of audits, the results of which are exhaustive. Whilst the 'flavour' of the paper is similar to several contributions to the 'optimal tax-evasion' literature, the enforcement procedure examined is, to the best of our knowledge, analytically distinct from all of these. It can, of course, be reduced to a special case of the general 'expected penalty' analysis of Becker (1968) and others, as (trivially) could any other model involving probabilistically levied penalties which has been written since 1968. As in all of these cases, however, the point of interest arises in examining the impact which different enforcement strategies have upon the characteristics of the expected penalty function, the deterrence structure thus generated and the impact of that deterrence structure upon the behaviour of the regulated agent.

The most important distinguishing characteristic of the tax-evasion context is that the taxpayer initially files a 'return' and it is on the basis of the level of income reported on this return that the agency decides whether or not it should conduct an audit. In these models (with the notable exception of Scotchmer and Slemrod (1989)) the audit, if conducted, is exhaustive. This class of models, however, does not incorporate the sort of uncertainty which is characteristic of models of non-compliance with environmental standards, and the problem facing the enforcement agency is comparable but distinct.

Section 1 provides a brief survey of the economic treatments of non-compliant behaviour in various contexts. It is far from exhaustive - many contributions, particularly to the literature on tax evasion, are very similar, and we survey only a representative sample. Chapter 2 of Roth, Scholz and Witte (1989) provides an insightful and thorough bibliographic survey.

In Section 2 we posit a simple, two-stage enforcement model. It is worthwhile in itself to analyse an enforcement set-up which has not been explicitly examined by previous authors. Our main motivation for the two-stage schema used in this paper, however, is that we believe it to resemble closely the enforcement régime actually employed by the EPA and other executive agencies in the domain of environmental protection in the United States. The firm's compliance problem is solved. Insofar as the characterisation of the enforcement environment is accurate, the model can be used to examine the impact of various proposed items of regulatory reform. Several results from the existing literature are duplicated.

In Section 3 the impact of changes in a policy parameter peculiar to a two-stage enforcement (the 'audit trigger') set-up are considered. An interesting distinction between the behaviours of a 'serious' and a 'non-serious' violator (in a sense to be defined) is identified. The distinction permits us to suggest an explanation for the co-existence of two 'stylised facts' of the US environmental scene in the 1980's - a coexistence which would be difficult to rationalise

In Section 4 we relax the assumption that monitoring uncertainty is necassarily exogeneous, argueing that it the firm will often have some discretion over its own 'inspectability' (transparency to inspection). We consider the ways in which firms may resist inspection, and their incentives to do so given the type of regulatory regime operated by the EPA and other major agencies of environmental protection. The 'conventional wisdom' that a non-compliant firm will necessarily benefit from an increase in monitoring uncertainty is disproved, whether or not the firm is permitted to re-adjust its rate of emissions ex post. The policy implications of endogenising the degree of monitoring uncertainty are outlined briefly.

(1) Selective Survey: Models of Non-Compliant Behaviour and Applications.

The so called 'Crime and Punishment' literature associated with, amongst others, Becker (1968), Stigler (1970) and Polinsky and Shavell (1979), has encouraged economists to analyse more thoroughly the issue of the enforcement of laws and the problems that arise when laws are not universally obeyed. Depending upon circumstances rules, laws and regulations are frequently vulnerable to being evaded, avoided or otherwise circumvented.

Analysis of incomplete enforcement has flourished in the public finance literature with the analysis of tax evasion. The taxpayer is assumed to weigh up the costs and benefits of evading the marginal dollar of tax that he pays. The taxpayer chooses to report some fraction of his actual income. There is some probability that his return will be audited, where that probability may itself depend in some way upon the characteristics of his report. The decision of how much tax to evade (ie not to report) depends upon factors such as probability of detection, size of fine for evasion, individual risk aversion etc, generally in 'intuitive' ways.

Early contributions treated the probability of an evader being caught as a fixed and uniform probability (the classic work is that of Allingham and Sandmo (1972)) and this remains a popular assumption (employed by, amongst many, Isachsen and Strom (1980), Langbein and Kerwin (1985), Alm (1988) and Tsebelis (1991)). Other authors have examined the scope for strategic auditing on the part of the collection agency. A more general literature exists on the optimal verification of contracts (see Baron and Besanko (1984) and the citations therein). The problem faced by the agency in this context is to determine for what values of reported income a return should be audited in order to maximise a specified objective function. Reinganum and Wilde (1986) find, for example, that when the penalty function is linear, for any given 'class' of taxpayers, those with greater true income under-report *less* than those with lower true income, and efforts at verification are lower the *greater* is reported income.

The compliance decision can be reformulated as a portfolio problem (Reinganum and Wilde (1986)). The agent must allocate his budget (monetary income) between a risky asset (namely unreported income) and a risk-free but lower yield asset (reported income). This permits all of the usual apparatus of portfolio theory to be brought to bear.

Other recent contributions have sought to consider either the impact upon the representative individual's compliance decision of alternative penalty or tax functions (eg Mookherjee and P'ng (1989, 1990), Clotfelter (1983), Crane and Nourzad (1985)), or of various changes in the detection process employed by the enforcement agency. Authors in this latter category have used increasingly sophisticated models of inspection in which the agency discriminates among taxpayers in various ways and according to various criteria. Particularly fruitful has been the employment of dynamic models in which the likelihood that a particular taxpayer will be audited in a given period depends upon that taxpayer's 'tax-history' (Landsberger and Meilijson (1982), Greenberg (1984)). The principal result is that, as is demonstrated in both of the papers cited, an enforcement régime in which the probability of inspection depends upon the outcome of the most recent inspection of a particular taxpayer is more cost effective than a system in which the taxpayer's inspection history is ignored.

Various aspects of the enforcement problem in the context of environmental-quality regulations have been analysed by Linder and McBride (1984), Harford (1978, 1987), Storey and McCabe (1980) and a number of other authors. The themes typically parallel those in the tax evasion literature. Though we will concentrate on those which deal with regulation based on emissions standards (that is where the regulator sets permissible limits on a firm's emissions rate and then penalises a firm caught exceeding those limits) we acknowledge the value of recent attempts to integrate partial compliance into emissions trading frameworks (notably the work of Malik (1984, 1990)).

Environmental economists have followed the public finance literature in formulating dynamic models in which inspection probabilities are state-dependent in one way or another. Several authors. Harrington (1988), Russell, Harrington and Vaughan (1986) and Russell (1984)) build upon Greenberg's (1984) tax model. They develop Markovian models in which firms are categorised according to their performance in past inspections. The probability that a firm will be inspected in the current period is dependent upon the category in which that firm currently finds itself. For example, a simplified version of the model would be one in which the probability that a given firm will be inspected depends upon whether or not it passed the most recent inspection that it faced. The movement of firms between categories is determined by a transition matrix. Two of the crucial differences between models of this type is the number of categories that they specify and how 'forgiveness' is modelled (how, that is, firms are able to get back into the 'good books' of the regulatory agency). One of the weaknesses of Greenberg's analysis is that once a firm has failed two consecutive inspections it finds itself in an absorbing state from which it cannot subsequently escape. If the possibility of type-1 errors is recognised (ie compliant firm will sometimes be ajudged to be non-compliant) then the model implies that in the long-run all firms will find themselves stuck in that absorbing state. It would be valuable to determine how far the outstanding welfare results of the model are robust to the relaxation of the absorbtive characteristic of the 'punishment' state. One recognition of particular significance in these dynamic models has been the possibility of penalty 'leverage'. Harrington (1988) and Harford and Harrington (1991) provide thorough discussions of this possibility in the environmental context.

The preponderance of contributions in the field employ static models, and the current paper is no exception. Various authors seek, in the context of static frameworks, to characterise the response of the firm to changes in various regulatory parameters.

Storey and McCabe (1980) construct a simple model of a 'criminal waste discharger' under emissions standards. They find that the actual rate of emissions by a firm decreases when the probability of inspection increases, fines become bigger and/or the emissions limit is reduced. All of these results are unambiguous and qualitatively intuitive. The model is limited, however, to the case in which the fine function is linear and the probability of detection is independent of the size of the violation. Storey and McCabe (1980: 32) predict that these restrictions save complexity without altering the major results. We demonstrate that a more realistic characterisation of the enforcement process will have important, qualitative impacts upon the responses of firms to marginal changes in policy parameters, and therefore upon the characteristics of the optimal regulatory régime.

One of the recurrent themes of the current paper is that different classes of firms will react to regulatory reform (as captured by changes in the regulatory parameters employed in our model) in quantitatively and, in some cases, qualitatively different ways. An important implication of this recognition is that optimal policy, where uniformly applied policy parameters must be set, is likely to be sensitive to the 'anatomy' of the regulated population - the proportion of the regulated population which can be fitted into each class.

Viscusi and Zeckhauser (1979) show that, when violations are punished by lump-sum penalties, reducing emissions limits below some critical point will cause compliance to fall; firms switch from being compliers to being violators. The size of a given violation is, however, independent of the standard. Adaire-Jones (1987) extends Viscusi and Zeckhauser's model by incorporating a generalised penalty function. She finds that if the marginal penalty function is non-decreasing then reduced emissions limits will cause increased compliance. If the marginal penalty is decreasing, however, tighter standards *may* cause *decreased* compliance. Harford (1978) also shows that tighter standards lead to greater compliance only when the marginal penalty is decreasing.

Kambhu (1991) develops a model in which the firm can expend resources to oppose penalties levied against it. He finds that tighter standards may generate increases in realised emissions regardless of the form of the marginal penalty function. As standards tighten the firm substitutes resources from investing in abatement effort to paying more and more fees to lawyers to oppose realised penalties. Kambhu's results are, however, built upon a rather simple-minded characterisation of the legal process. In particular, letting p be the size of the penalty that the firm should pay in some given case, and h be the money spent by the firm on lawyers to oppose the fine, he assumes that $\partial p/\partial h$ is everywhere negative.

(2) A Two-Stage Characterisation of the Regulatory Monitoring Process.

A firm chooses to produce e units of 'emissions' (pollution) in going about its business. It's total costs of production are $c(e)>0$, where $\partial c/\partial e<0$ and $\partial^2 c/\partial e^2>0$. (From now on subscripts and primes will denote derivatives in the usual way). It is cheaper to operate using a 'dirty' technology.

We consider a stylised two-stage enforcement process under uncertainty. The regulator sets a standard E such that a firm is considered to be in violation when $e>E$.

The regulatory bureau then conducts two classes of monitoring excercise, which we term 'inspections' and 'audits'. Inspections are preliminary, non-exhaustive assessments of the compliance status of the firm and yield 'preliminary evidence'. The taking of inspections is random: in a given period the firm will be inspected with probability $\mu>0$. An inspection yields an estimate of e, ê. We define the distribution F(ê|e,s) as F((ê-e)/s) (where s is a positive scale parameter). The associated density, f((ê-e)/s), is symetric, everywhere positive, and increasing (decreasing) in its argument when ê<e (ê>e). The parameter s measures the degree of dispersion of ê around e and will be interpreted as the 'noisiness' of the inspection process.

If the evidence collected at the inspection is sufficient to suggest that the firm is indeed in violation (ie ê is 'sufficiently' large compared to E) then the regulator will conduct an environmental 'audit' - continued investigation into the firm with a view to constructing a case for civil or criminal prosecution. An 'audit' is triggered when ê exceeds some critical level Ω, ie if ê>Ω. We assume that this second round of 'non-routine' investigation is exhaustive in that it will yield the truth about the firm's compliance status. The 'audit trigger', Ω, is a regulatory parameter over which the agency has discretion. It parameterises the referral practices of the regulatory agency. When Ω is large the referral practices are relatively lax in the sense that the preliminary evidence collected by the initial round of inspection has to be relatively conclusive in order to induce the agency to conduct an audit and push for prosecution of the firm in question.

It is on the basis of audit results that the firm is penalised, according to the penalty function $\emptyset(e,E)$. It is reasonable to suppose that $\emptyset=0$ for $e\leq E$, $\emptyset>0$ elsewhere. We restrict ourselves to the class of penalty functions where the only determinant of the penalty is the magnitude of the violation, if any (ie those that could be respecified in the form $\emptyset(e,E) = \S(e-E)$). The most significant implication of this is that $\emptyset_e=-\emptyset_E$. Throughout the paper it will be assumed that the penalty is a non-decreasing function of the magnitude of the violation (that is, $\emptyset_e\geq 0$), but no restriction is put upon the sign of $\emptyset_{ee}$ which will characterise whether the penalty function is progressive, neutral or regressive.

With the firm choosing to emit e units, the probability that an inspection at the firm will lead to an audit will be denote $a=a(e,\Omega,s)$. It is clear that

$$(1) \quad a(e,\Omega,s) = 1-F((\Omega-e)/s)$$

For analytical as well as interpretive convenience we will distinguish between two classes of non-compliant firms. The 'non-serious violator' (NSV) is defined as a firm for which $E<e<\Omega$. For the 'serious violator' (SV) $e>E$ and $e\geq\Omega$. We demonstrate in an appendix that $a_e>0$, $a_\Omega<0$, that for a NSV $a_{ee}>0$, $a_{\Omega\Omega}>0$, $a_{e\Omega}<0$ and $a_s>0$, whilst for a SV $a_{ee}<0$, $a_{\Omega\Omega}<0$, $a_{e\Omega}>0$ and $a_s<0$.

The qualitative differences in some of the signs generates the analytical significance of the distinction between the two classes of non-compliant firms. For example, a_e measures the impact that a marginal increase in the rate of emissions by the firm has upon the probability that if the firm is inspected an audit will be triggered. If $e>\Omega$ (ie we are dealing with an SV) then this marginal impact decreases as e increases (or, equivalently, as Ω decreases) because Ω is sliding <u>down</u> the left hand side of the density function. That is $(\partial^2 a/\partial e^2)=-(\partial^2 a/\partial e \partial \Omega)<0$ when evaluated for the serious violator. The inequality is reversed in the case of a non-serious violator. This is because, atleast whilst e remains less than Ω, an increase in e slides Ω up the right-hand tail of the density function.

In later sections we investigate how this and other qualitative differences between SV's and NSV's will affect the firms 'pollution decision' at the margin. How a non-compliant firm will react to changes in various changes in regulatory parameters is found to depend upon the initial degree of non-compliance. It is misleading to distinguish only between 'compliers' and 'non-compliers', as the existing literature does.

(2.1) Motivation For the Regulatory Scenario.

Our two-stage characterisation and the meaning of Ω can be justified in at least two ways, and is motivated by the practices of the Environmental Protection Agency (EPA) in the US (see Mintz

(1989) and Strock (1990) for more detailed institutional accounts of the practices of the USEPA. It is similar to the 'basic investigation model' of Baiman and Demski (1981).

First, Ω can be viewed as an institutional constraint upon the agency. Thus EPA inspectors are able to make supervised visits to plants by appointment or can monitor a firm's compliance from outside of its premises (eg by air sampling). The firm's rights to privacy under the 4th amendment of the US Constitution require that the agency convince a judge that there is 'reasonable cause to believe that an offence is being or has been committed' in order to be granted a search warrant. Thus Ω could be regarded as the judicial interpretation of what constitutes 'reasonable evidence' that $e>E$. If the judge grants a warrant (ie if $\hat{e}>\Omega$) then the EPA refers the case to its 'National Enforcements Investigations Center'. The NEIC is staffed by fully deputized US marshals who enter the offending plant without notice to execute the search warrant, seize relevant evidence and subpoena plant operatives as required. These second-round investigations are what we refer to as 'audits' and can be regarded as exhaustive.

According to this view the agency is assumed willing and able to audit a regulated firm whenever the court gives it permission to do so. The court's interpretation of the rights of the regulatee and the requirements of due process constitute the binding constraint upon when an audit will and will not take place.

Alternatively, Ω can be interpreted in a more general context as an inevitable consequence of a regulatory agency facing a binding budget constraint. If the regulator does not have the manpower to audit every firm then he is obliged to 'prioritise' - to order audits of only those plants which preliminary evidence from inspection suggests are the most likely to be found to be in violation of standards. This will translate into some trigger value Ω which a firm must satisfy in order to avoid being amongst those which the regulator decides to audit. The regulator is assumed able, through whatever means, to be able to commit to the values of his policy instruments and is.

According to this second view of the world, the constraint of due process which the court imposes is non-binding. There will be a range of values for $\hat{e}$ to which the agency will not respond with an audit even though the court would be willing to authorise one.

In the sections which follow we will examine the behaviour of a representative firm. We employ, except where otherwise stated, the (empirically justifiable) 'small-firm' assumption, namely that the individual firm is so small relative to the industry that it can treat the regulatory parameters (Ω, E and μ) as predetermined.

(2.2) The Firm's Problem.

The firm chooses its level of emissions to minimise a loss function, L, the sum of compliance costs and regulatory penalties. The firm's problem and the necessary conditions for an interior solution can be represented as follows.

(2) $\text{Minimise}_e\ L : \{c(e) + \mu.a(e,\Omega,s).\text{Ø}(e,E)\}$

(3) $\partial L/\partial e = 0 = c'(e) + \mu a_e \text{Ø} + \mu a \text{Ø}_e$

(4) $\partial^2 L/\partial e^2 = c''(e) + \mu.\left[a_{ee}\text{Ø} + 2.a_e\text{Ø}_e + a\text{Ø}_{ee} \right]$

The second-order condition $(\partial^2 L/\partial e^2)>0$ is assumed to hold over the relevant ranges (that is, the stationary point defined by the solution to Equation 3 is assumed to be a minimum).

(2.3) The Compliance Case.

If the required standards are sufficiently lax (i.e., if E is sufficiently large) then the firm will choose to comply with the standard. In effect the corner solution dominates all interior solutions. At some point, as requirements are tightened (as E is reduced), the firm chooses to violate.

Notice that the firm will never choose to over-comply. This is because c(e) is indefinitely increasing in e. The assumption that the audit (as opposed to the inspection) technology is perfect means that there is no chance of a compliant firm being wrongly penalised due to inaccuracies in the audit technology (though the inspection process is imperfect and some firms which are referred for audit may subsequently be found to be in fact compliant). Thus, having complied, there is no need for the firm to go any further.

More concretely, the firm will choose to comply (ie to choose $e=E$ if and only if $E>E_c$ (E_c can be thought of as the critical level of regulation at which the firm is indifferent between compliance and optimal non-compliance) where E_c satisfies;

$$(5)\quad c'(E_c) + \mu.a(E_c,\Omega,s)\emptyset_e(E_c,E_c) = 0$$

When $E=E_c$, the firm's first-order condition is satisfied exactly when $e=E$. At $E=E_c$ the reduction in production costs associated with marginal non-compliance are exactly offset by the increase in expected penalties simultaneously generated. The impact that changes in the principal parameters of the firm's problem have upon its decision to comply (ie upon E_c, how tight standards have to be in order for the firm to be induced to break them). If firms are distributed according to marginal cost characteristics ($\partial c/\partial e$) then an increase in E_c, for any given E, can be interpreted as an increased likelihood that a randomly chosen firm will be a non-complier.

Proposition (1): The critical level of the environmental standard at which the firm decides not to comply is a decreasing function of the probability of inspection and the marginal penalty in the vicinity of $e=E$, and an increasing function of the audit trigger (Ω). It may be an increasing or decreasing function of the accuracy of the monitoring technology employed by the inspectorate, depending upon the initial compliance status of the firm.

Proof: Let the fine levied upon a firm shown to be in violation of the emissions standard by some arbitarily small epsilon be denoted t. Application of the implicit function theorem to the equality characterising E_c (i.e. the F.O.C. evaluated at e=E) yields the following derivatives;

$$(6)\ \partial E_c/\partial \mu = -\left\{ a(E_c,\Omega,s).t \Big/ \left[c''(E_c) + \mu.a_e(E_c,\Omega,s).t + \mu.a(E_c,\Omega,s).\emptyset_{ee}(0) + \mu.a.\emptyset_{eE}(0)\right]\right\}$$

Recalling that $\emptyset_{ee}(0)=-\emptyset_{eE}(0)$ (a simplification that we will use in each of the following cases), this can be simplified to

$$(7)\ \partial E_c/\partial \mu = -\left\{ a(E_c,\Omega,s).t \Big/ \left[c''(E_c) + \mu.a_e(E_c,\Omega,s).t\right.\right\} < 0$$

similarly,

$$(8)\ \partial E_c/\partial t = -\left\{ \mu a(E_c,\Omega,s) \Big/ \left[c''(E_c) + \mu a_e(E_c,\Omega,s).t\right]\right\} < 0$$

$$(9)\quad \partial E_c/\partial\Omega = -\left\{ \mu a_\Omega(E_c,\Omega,s).\varnothing_e(E_c,E_c) \Big/ [c''(E_c) + \mu a_e(E_c,\Omega,s).t] \right\} > 0$$

$$(10)\quad \partial E_c/\partial s = -\left\{ \mu a_s(E_c,\Omega,s).\varnothing_e(E_c,E_c) \Big/ [c''(E_c) + \mu a_e(E_c,\Omega,s).t] \right\}$$

The signs are as indicated. In the final case, $\partial E_C/\partial s$ is positive and negative for a SV and NSV respectively. *QED.*

Thus the firm will comply with a comparatively stringent standard (ie E_c is low) when inspections are frequent (μ is large), when the marginal penalty at e=E is comparatively large and/or when referral practices are stringent (ie the value of the audit trigger, Ω, is small). The absolute size of each of these effects is negatively related to $c''(e)$, the rate of increase of the marginal cost of emissions reduction.

The ambiguity in the sign of $\partial E_c/\partial s$ is of some interest. Recall that s is a scale parameter capturing the degree of dispersion of ê around e. A decrease in s can be interpreted as either an improvement in the monitoring technology employed by the regulatory inspectors or an increase in the 'transparency' of the production technology of the firm (in the sense of it being 'easier' to assess the emissions resulting from the production process given a monitoring technology of unchanged sophistication). When s is smaller $\hat{e}$ is a less noisy estimator of e.

'Conventional wisdom' suggests that $\partial E_c/\partial s$ would be positive. In that case an increase in the transparency of the technology employed by the firm will serve to reduce E_c - to increase the critical stringency in standard at which point the firm finds it optimal to violate. A necessary and sufficient condition for this is that $a_s(E,\Omega,s)$ be negative which, in turn, is necessarily the case when $\Omega<E$. If $\Omega>E$, so that the referral rule employed by the regulator is sufficiently tight vis-à-vis the standard, then the sign of the derivative is reversed. In this case an increase in the transparency of the firm's production technology would increase the range of values for the standard at which the firm would choose not to comply.

The possibility that $\partial E_c/\partial s$ may be negative seems perverse but is likely to be the case in the plausible circumstances where the referral practices of the regulator are such that the standard is exceeded by the audit trigger, $\Omega>E$. This condition holds if, for instance, the emissions limit is 100 units per firm and the regulator conducts an audit on any firm where preliminary emissions readings exceed 110 units. The qualitative ambiguity in the role of inspection uncertainty in the compliance decision can be understood by considering the impact of changes in s upon the probability of audit *conditional on having been inspected.*

An increase in the transparency of the firm's technology, a decrease in the scaling parameter s, is associated with a reduction in the dispersion of the probability density around e. For a non-serious violator (NSV) a reduction in the parameter s serves to decrease the conditional probability that an inspection will trigger an audit. Conversely, for a serious violator (SV), with $\Omega<E$, the same reduction in the scaling parameter s leads to an increase in the analogous probability.

Increases in the transparency of the technology used by a firm, in the sense that an inspector can assess the firm's true emissions level more accurately, may, ceteris paribus, increase the range of values for the standard over which the firm will choose not to comply. The determining factor will be the stringency of the regulatory bureau's referral practices.

(2.4) Optimal Non-Compliance.

The behaviour of a firm that exceeds the standard is of great empirical interest, since the rate of compliance with a number of environmental regulations in the US and abroad is extremely low. For policy analysis to be relevant it must take into account the behaviour of the non-compliant as well as the compliant firm and the interaction between policy parameters and the non-compliance decision.

There are four forms which 'stricter' regulation might take in the construct presented here (though later we will argue that only those marked (ii) and (iii) are under the ongoing discretion of the enforcement agency);

(i) Tightened emissions standards (ie reductions in E).
(ii) Tightened referral practices (ie reductions in Ω).
(iii) Increased inspection intensity (ie increases in μ).
(iv) Increased fines for violation (ie increases in Ø).

Below, we examine the impact of changes in each of these parameters, though it will be argued that not all of them will, in general, be under the day to day control of the regulatory agency. The level of fines, for example, will often be determined by sentencing courts in line with criteria other than technical efficiency. It is demonstrated that only increased inspection intensity and increased fines will unambiguously generate an improvement in environmental quality when the firm is not initially in compliance.

To examine the interior solution case we recall the first- and second-order conditions (Equations 3 and 4 respectively) associated with the firm's problem (where asterisks now denote solution values);

$$(11)\quad \partial L/\partial e = 0 = c'(e^*) + \mu a_e(e^*,\Omega,s).\emptyset(e^*,E) + \mu.a(e^*,\Omega,s).\emptyset_e(e^*,E)$$

$$(12)\quad \partial^2 L/\partial e^2 = J(e^*,E,\Omega,s) > 0$$

The first-order condition dicates that the firm will increase emissions up to the point at which the marginal cost-savings from using 'dirtier' technology balance the associated increases in the expected regulatory penalties which follow from producing more emissions. This is a simple analogue to a result common in the 'crime and punishment' literature (associated with the classic articles by Becker (1967) and Stigler (1972)). In this context the expected penalty increase has two components - an increase in the likelihood that an inspection, should one be conducted, will yield enough preliminary evidence to trigger an audit, and an increase in the penalty faced if the the true size of the firm's violation is indeed uncovered.

As above, the second partial derivative of the loss (objective) function with respect to e will be represented by 'J' for convenience: $J=J(e,E,\Omega,s)$, though the arguments will generally be supressed for brevity. For e^* to constitute a minimum it is necessary to assume that $J(e^*,E,\Omega,s)>0$.

Given (empirically) the preponderance of non-compliant behaviour, it is of interest to consider the factors which affect the firm's choice of the severity of its violation. In interpreting the propositions in the next section it should be noted that the population of firms is likely to be heterogeneous - consisting of compliers, non-serious violators and serious violators in some proportions. The impact of the regulatory reform considered upon aggregate emissions will, then, be some sort of weighted average of the firm level impacts.

(2.5) Tightening Emissions Limits and Non-Compliant Behaviour.

Proposition (2): Reduction in the emissions limit may induce an increase or a decrease in the emissions output of a non-compliant firm.

Proof: Simple application of the implicit function theorem to the non-compliant firm's FOC (Equation 11) yields

$$(13)\quad \partial e/\partial E = -\mu\left\{[a\O_{eE} + a_e\O_E] / J\right\}$$

$$= (\mu/J).(a_e\O_e + a\O_{ee})$$

$$= (\mu/J).a.\O_e.[(a_e/a) + (\O_{ee}/\O_e)]$$

J is positive by the second-order condition for a minimum. The sign of $\partial e/\partial E$ is the same as the sign of the numerator. The sign of the numerator and thence the sign of $\partial e/\partial E$ is ambiguous for both types of violators. *QED*

Proposition (3): An increasing marginal penalty function is a sufficient but not necessary condition to ensure that increasedly stringent standards will induce a reduction in emissions by a non-compliant firm.

Proof: Consider the expression derived for $\partial e/\partial E$ in proposition (2). It is straightforward to see that $\partial e/\partial E>0$ if

$$(14)\quad [a_e Ø_e + aØ_{ee}]>0$$

which is necessarily satisfied when $Ø_{ee}>0$, but is also satisfied when $Ø_{ee}<0$ provided $a_e Ø_e$ is sufficiently large. *QED*

Proposition (4): If violations are punished according to a linear penalty schedule then an increased stringency of emissions standards will, ceteris paribus, induce a decrease in emissions by all violators.

Proof: In this case $Ø_e>0$, $Ø_{ee}=0$. An increased stringency of standards is associated with a decrease in E. The result follows directly from Proposition (3). *QED*

Notice that a regressive penalty function (i.e. $Ø_{ee}<0$) makes it less likely that the representative firm's problem will have an interior solution (see Equation 5 which chsaracterises E_c). It does not, however, rule out the possibility - being neither a necessary nor sufficient condition for compliance to be the firm's dominant strategy. In this section we are conducting comparative statics on the solution which is assumed to be interior. If the solution is not interior then it is not characterised by non-compliance, and the analysis of it is of limited interest to us.

These simple results characterise the impact that changes in emissions standards can be expected to have upon the emissions of compliant and non-compliant firms. They are extensions of the results of Adaire-Jones (1987) and Harford (1978), outlined in Section (1), to the case of a two-stage enforcement régime. They contradict the result of Storey and McCabe (1980), which the authors claim to be robust to changes in the form of the penalty function and the manner in which the probability of detection varies with size of violation (Storey and McCabe (1980: 32)), that $(\partial e/\partial E)$ will be everywhere positive for non-compliant firms. Viscusi and Zeckhauser (1979) examine the case of lump-sum penalties and find that $(\partial e/\partial E)>0$ for any individual non-compliant firm - as in Proposition (4) presented here.

Of more interest is the demonstration that the tightening of standards may reduce environmental quality by inducing an increase in the equilibrium rate of actual emissions. Furthermore, $\partial e/\partial E$ may be qualitatively different for different violators. The result contrasts with Kambhu's assertion that "...relaxing a regulatory standard will *always* produce an increase in performance by non-complying firms", (Kambhu (1991: 109). It provides analytic support for the conclusions that Craswell and Calfee draw from a series of simulation excercises that "....where the uncertainty is distributed normally.....some relaxation of the legal standard *might* improve a defendant's compliance decisions", (Craswell and Calfee (1986: 299)) only where the marginal penalty function is regressive.

(2.6) Increasing the Probability of Initial Inspection.

It is generally supposed that one of the most important determinants of the frequency of non-compliant behaviour is the probability of inspection. If a plant is not inspected then any contravention it may have committed will be unpunished.

Proposition (5): An increase in the probability of initial inspection induces, ceteris paribus, a decrease in the equilibrium level of emissions and the equilibrium absolute size of violation of both the serious violator and the non-serious violator.

Proof: Implicit differentiation of Equation 11 yields;

$$(15)\quad \partial e/\partial\mu = \partial(e\text{-}E)/\partial\mu = -\left\{[a_e.\text{Ø} + a.\text{Ø}_e]/J\right\} < 0$$

The denominator of the term in braces is everywhere positive by the soc for a minimum. Both constituents of the numerator of the term in braces are positive in the both the SV and the NSV case. The sign of $\partial e/\partial \mu$ is therefore unambiguously negative when evaluated for either the SV or the NSV. Since E is constant by hypothesis, $(\partial(e-E)/\partial \mu)=(\partial e/\partial \mu)$ and is also unambiguously negative for both classes of violator. *QED*

Thus, by increasing the frequency with which it carries out routine inspections the regulatory bureau is unambiguously able to achieve improvements in environmental quality by inducing reduced emissions from both classes of violator. The relative emissions reductions achieved by the SV vis-a-vis the NSV (ie the relative magnitude of $\partial e(SV)/\partial \mu$ vis-a-vis $\partial e(NSV)/\partial \mu$) cannot be determined without knowledge of various third derivative terms. The result of Proposition 5 supports those of Storey and McCabe (1980) and Harford (1978). In a later paper Harford posits a model in which the firm self-reports its rate of emissions, and can be penalised for misreporting, as well as for exceeding the emissions limit *per se*. (Harford (1987)). Although he determines that $\partial e/\partial \mu$ will be non-positive, he is able to identify plausible scenarios which will ensure that $\partial e/\partial \mu=0$. In the current paper, without self-reporting, the inequality is strict.

It is interesting to notice the qualitatively different roles of increased inspection intensity versus increased audit intensity. Account should be taken of this relationship in allocating regulatory funds within the agency. The budgetary problem is not simply one of determining how much money to expend on 'enforcement' in aggregate, but also to decide how best to divide it between spending on inspection and audit effort.

(2.7) The Role of Increased Penalties.

The level of penalties is one of the two parameters in our model which could be changed, at least in theory, by fiat (the other being the emissions limit).

In our analysis we regard the emissions limit and the penalty schedule as being outside of the control of the regulatory agency. This is a plausible supposition. The former is mandated legislatively at infrequent intervals, whilst the latter is the jurisdiction of the civil and criminal courts. Courts can generally be seen as administering punishments which are seen as 'fair and proper' in accordance with established sentencing practices.

Whilst the regulator may seek to impress upon the judiciary the seriousness of environmental damage, in the hope that courts will reassess what they regard as 'proper' retribution in the case of infringements of environmental legislature, the penalty function ($\emptyset(e,E)$ in our terminology) cannot be regarded as policy manipulable in the usual sense, though it may be possible (in a dynamic context) for the U.S.E.P.A. to apply what is called 'leverage'. By making the probability that the firm will face inspection in future periods dependent upon compliance performance in the current period the regulator can ensure that the expected value to the firm of being caught in violation in this period is greater than the nominal dollar value of the (often modest) fine levied. The future loss of surplus to the firm by having been cast into a pool to be faced with a harsher future environmental régime must be added to the fine itself in order to generate an assessment of the 'effective fine'. In this sense the regulator is said to apply 'leverage' and, as such, to excercise discretion over the effective penalty function even when he doesn't have discretion in the determination of fines. Penalties assessed by courts will reflect some exogeneous concept of justice (in the sense of just retribution given the gravity of the offence) rather than concerns regarding the place of penalties in a policy-makers calculation of optimal policy.

Penalties faced by non-compliant firms in the US have been notoriously small. Table 1 presents some evidence of this, including calculations of the average penalty assessed and the average penalty per 'notice of violation' issued. This latter value can be regarded as a proxy for the expected penalty per violation uncovered. It can be seen that whilst it varies greatly according to the state in question the dollar values are have been very low. They have tended to increase, however, in recent years. Moreover, courts have become more ready to supplement monetary civil penalties, of the type usually envisaged in the case in which a corporation fails to comply with some requirement, with criminal prosecution of responsible corporate executives. Selected data on this trend is tabulated in Table 2 (see, for detailed description of cases, Mealey's Litigation Reports (December, 1988)). In 1988, for example, the US Department of Justice (on behalf of the EPA) prosecuted 97 defendants in such cases - a seven-fold increase

State	NOV's Issued	Average Penalty Assessed	Average Penalty Per NOV
Colorado	124	120	0.5
Kentucky	194	2520	68
Minnesota	41	10,900	266
Nevada	32	45	3.3
Virginia	161	200	3.8
Wisconsin	80.5	7951	760

Table 1: Enforcement Statistics for Selected States (annual averages, 1978-83). Columns show, from left, total number of 'Notice of Violations' issues, Average dollar value of penalty when assessed and average penalty per Notice of Violation. (Source: adapted from Russell (1990: 254)).

	1982	1985	1988
Cases Prosecuted	7	15	27
Defendants Charged	14	40	97
Days Sentence per Convicted Defendant	-	78	278
Months Probation per Convicted Defendant	19	22	26

Table 2: EPA Criminal Enforcement Activities, Selected Statistics 1982-1985. (Source: Russell (1990:267) and Mealeys Litigation Report (1988, 1(18)).

on 1982. It is also noticeable that the sentences upon conviction have become harsher, averaging eight months for each of the 58 defendants convicted in 1987. Whilst the criminal prosecution programme remains small and has not yet run for long enough for it's impact upon industry to be assessed adequately, it may be that the ultimate threat of criminal detention may prove a strong deterrent in the fight against environmental crime. In particular it may seek to deter the 'calculating violator' who has, up to now, been able to regard monetary environmental penalties assessed against his firm as just another entry in the accounts - as undesireable as any other expense but just another 'cost of doing business'.

A number of observers have commented that it would be of great policy value to the EPA to be given increased discretion over penalties assessed for infringements of regulations under its jurisdiction. Analytically such a move would make $\emptyset(e,E)$ a choice in the characterisation of optimal policy - giving the policymaker a fiat policy instrument. The legal ramifications of such a move, however, make it problematic.

It is still of interest to specify how changes in the level of penalties might affect the comportment of regulated firms. Consider a parallel upward displacement of that part of the penalty schedule to the right of E, leaving the structure of marginal penalties unaltered except at e=E (at which point is plausible to suppose that the penalty function may exhibit a discontinuity)

In this case the penalty can be thought of as analogous to a 'two-part tariff' - a lump-sum is payable even by an 'epsilon-violator' (that is a violator who is in excess of the emission limit by an arbitrarily small amount), and the penalty beyond that is increasing in the seriousness of the violation. This would seem to be a reasonable stylisation of the EPA's 'two-part' view of penalty assessments.

Proposition (6): An increase in the level of penalties, keeping the marginal penalty structure unchanged over relevant domains, induces a decrease in the equilibrium level of emissions and in the absolute magnitude of violation by both a serious and a non-serious violator.

Proof: Suppose that $\emptyset(e,E)=\alpha+\varphi(e-E)$, for $e>E$. Application of the implicit function theorem to the non-compliant firm's FOC yields

$$(16)\quad \partial e/\partial\alpha = \partial(e-E)/\partial\alpha = -\left\{\mu . a_e/J\right\} < 0.$$

The denominator is everywhere positive by the soc for a minimum. The numerator of the term in braces is positive for both the SV and the NSV. Thus $\partial e/\partial\alpha$ is negative for both classes of violator. *QED*

Thus the absolute *level* of fines does affect the decisions of the firm. This is contrary to the assertion often made that it is only the *marginal* characteristics of the penalty function which affect the solution to the firm's problem.

(3) Two-Stage Enforcement, Tightening Referral Practices and Non-Compliant Behaviour.

Whereas control of E, the emissions standards, is not under the day-to-day control of the regulatory agency, the agency does have discretion (subject to budgetary constraints) of altering Ω. Tightening of referral practices is a commonly observed method by which agencies try to effect improved compliance. The impact of this type of regulatory reform in a two-stage enforcement context is summarised in the following proposition;

Proposition (7): A tightening of referral practices will, in equilibrium, induce a decrease in the emissions level and absolute size of violation of the non-serious violator. It may induce an increase or decrease in the same variables for the serious violator.

Proof: A tightening of referral practices is associated with a decrease in the audit trigger, Ω. Application of the implicit function theorem to the foc of the firms problem yields an expression for $\partial e/\partial\Omega$. Attaching signs to the various constituents for the case of the NSV and SV gives, respectively;

$$(17) \quad \partial e/\partial\Omega(\text{NSV}) = -\mu.\left\{[a_{e\Omega}(e_{NSV},\Omega,s).\emptyset(e_{NSV},E) + a_{\Omega}(e_{NSV},\Omega,s).\emptyset_e(e_{NSV},E)] \Big/ J\right\} > 0$$

$$(18) \quad \partial e/\partial\Omega(\text{SV}) = -\mu.\left\{[a_{e\Omega}(e_{SV},\Omega,s)\emptyset(e_{SV},E) + a_{\Omega}(e_{SV},\Omega,s).\emptyset_e(e_{SV},E)] \Big/ J\right\}$$

Thus, $\partial e/\partial\Omega(\text{NSV})$ is unambiguously positive. The sign of $\partial e/\partial\Omega(\text{SV})$ is ambiguous and is the same as the sign of the numerator in the expression given above. With E fixed $\partial(e-E)/\partial\Omega=\partial e/\partial\Omega$ in each case, i.e. induced changes in the equilibrium level of emissions are associated with one-to-one changes in the absolute size of the firm's equilibrium violation. *QED*

The difference in result between the serious and non-serious violator follows from the difference in the sign of $a_{e\Omega}$ between the two cases. Unlike in the case of the non-serious violator, for the serious violator an increase in Ω serves to reduce the additional increase in the risk of being audited induced by the incremental unit of pollution ($a_{e\Omega}$ is negative for the NSV).

This is an interesting result. If environmental standards are fixed (by legislative mandate, for instance) such that

$$(19) \quad a_{e\Omega}(e_{SV},\Omega,s).\emptyset(e_{SV},E) > -a_{\Omega}(e_{SV},\Omega,s)\emptyset_e(e_{SV},E),$$

then a decision by the bureau to tighten referral practices by reducing the audit trigger (i.e., reducing the 'amount' of preliminary evidence that the first round of inspection must yield in order for the inspected firm to face an audit) will induce an *increase* in emissions by serious violators. This type of apparently perverse reaction by the firm to the regulatory change is particularly likely when the level of the penalty faced by a serious violator who is convicted, $\emptyset(SV)$, is large relative to the marginal penalty $\emptyset_e(SV)$. More formally, a sufficient condition for $\partial e/\partial\Omega$ to be negative (ie the 'perverse' case) is that the elasticity of the marginal audit probability with respect to changes in emissions be sufficiently large vis-a-vis the elasticity of the penalty function with respect to increased emissions. In other words, after some rearrangement, the above condition reduces to;

$$(20)\quad (\partial\emptyset/\partial e).(e/\emptyset) < (\partial a_\Omega/\partial e).(e/a_\Omega)$$

In cases where such a condition does hold (where courts hand-out sufficiently punitive penalties) the regulator can encourage decreased emissions from firms by *relaxing* the referral practices, that is, by being less aggressive in ordering audits by raising the audit trigger, Ω.

The source of the apparent perversity can be seen by inspecting the terms in the expressions for $\partial e/\partial\Omega$ in Proposition 7. The optimising firm equates the 'full' cost of the marginal unit of emission against its 'full' benefit. In the case of a NSV a tightening of referral practices (ie a fall in Ω) has two impacts. It increases the impact which the marginal unit of e has upon the probability that the firm will be audited if it is inspected, and thus be liable for the fine $\emptyset$. This effect makes the marginal unit of emission more 'expensive', particularly when $\emptyset$ is large. The fall in Ω also increases the conditional probability of audit directly. In so far as the marginal unit of e has increased the penalty payable in this case, the marginal unit of emission has again become more 'expensive'. The two effects are captured by the first and second composite elements of the numerator of the term in braces, both of which are negative. In the case of the NSV both effects work to make the marginal unit of emission less attractive for the firm to carry on generating and thus it chooses to reduce e. In the case of the SV, as is shown by the signs attached to the two relevant terms in the proposition, the two effects work against each other; hence the ambiguity in sign.

The possibility that $\partial e/\partial\Omega$ will be negative for some firms, so that a tightening of the audit referral practices will cause those firms to increase the degeree to which they violate environmental standards, is not merely a mathematical possibility but quite plausible in many real-world contexts. It is merely required that the amount of the fine be sufficiently large vis-à-vis the 'steepness' of the fine function. A sufficient but not necessary condition for this requirement to be satisfied is that penalties be lump-sum, at least over the relevant range. A two-part penalty function characterised by a stiff lump-sum and a gently sloping fine function may also suffice.

Ambiguity regarding the impact of changes in Ω or E upon the emissions behaviour of violators of different severity make it prudent that the regulator consider the efficacity of alternative policy instruments.

(3.1) The 'Anatomy' of the Regulated Population and the Choice Between Inspections and Audit Effort.

The aim of this paper is to examine the appropriate choice of policy instrument to achieve environmental improvement in the plausible context in which monitoring technology is imperfect and compliance is less than full. The results highlight the dangers of basing policy discussion on various apparently 'obvious' assumptions.

For example, according to Proposition 2 only in certain cases will *both* $\partial e(SV)/\partial E$ and $\partial e(NSV)/\partial E$ be positive. If one or the other of the derivatives is negative then tightening emissions limits will have qualitatively different impacts upon the equilibrium rate of emissions of firms which initially violate the existing standard with different degrees of severity.

More ambiguities apply in the impact of tightening of referral practices. In this case similar comments apply. Proposition 7 demonstrates that there may be qualitative differences in the effect of tightened referral practices on rates of emission ($\partial e(NSV)/\partial \Omega$ is positive but the sign of $\partial e(SV)/\partial \Omega$ is ambiguous). This is particularly likely when the level of penalties is high.

Propositions 5 and 6 demonstrate that the effect of increased inspection (as opposed to audit) probability and increased level of penalties is unambiguously to reduce the equilibrium rate of emissions for both classes of violator. They are, in this sense, 'fool-proof' policy instruments. It should be recognised, however, that increasing the inspection intensity is likely to be a costly activity for the bureau and thus less preferable on those grounds alone. In addition, as was discussed in Section 2, the level of penalties is not generally at the discretion of the regulatory agency (where violations are prosecuted through civil or criminal courts) but is determined by established guidelines and precedents for sentencing.

The recognition that firms which differ according to their initial compliance status (that is, as members of the mutually exclusive and exhaustive classes, 'compliers', 'non-serious violators' and 'serious violators') will react to policy reform in qualitatively as well as quantitatively differing ways is an important one in policy discussion. In general we have supposed that $E<\Omega$. This is the richest case and likely to be the most empirically plausible in most contexts (in the case of the EPA, for example). If $\Omega<E$ then all violations necessarily constitute 'serious' violations (in the sense of this paper) and the analysis is simplified. A reform which may cause a typical non-serious

violator to 'clean up his act' may induce a serious violator to do quite the opposite. The impact upon aggregate emissions of any particular policy reform, then, will depend crucially upon the 'anatomy' of the regulated population. Furthermore, distinguishing between compliers and violators is inadequate. There are significant qualitative differences in reactions among violators.

Throughout the paper it has been supposed that the regulatory agency is restricted to formulating 'uniform' policy - policy which applies uniformly to all members of the regulated population. This is the assumption generally made by scholars in this area of the economic literature on regulation. If it is supposed (quite plausibly) that the cost function of the firm is not known by the regulator then the uniformity assumption follows naturally from the decision to employ a static model. When working with a static model, however, we should not lose sight of the fact that the real-world phenomenon which we seek to understand is an essentially dynamic one, involving repeated interaction between regulator and firm in an ongoing economic environment. In such situations it may be possible for the regulator to 'learn' things about the firm and to apply them later. Through such means different rules can be applied to different firms according to the histories of their respective interactions with the regulatory agency. The gains from non-uniform regulation have been examined by a number of authors (see Harrington (1988) and Harford and Harrington (1991)) and may be substantial. Analogous analysis has been presented, as was mentioned in Section 1, by Greenberg (1984) and others in the context of the problem faced by the revenue service in ensuring taxpayer compliance.

For the purposes of the present analysis we restrict our attention to uniformly applied policy, whilst noting the potential inadequacies of so doing. The two policy instruments over which we assume that the regulatory agency excercises discretion are the inspection intensity (denoted μ), and the stringency of referral practices (Ω). Recall that μ is merely the probability that a given source will be visited by an inspection team and, as such, lies on the closed interval [0,1]. As such an increase in μ constitutes an increase in inspection intensity. The parameter Ω designates the 'amount' of preliminary evidence ($\hat{e}$) required to trigger a second round of investigation - what we have termed an audit. An increase in Ω, then, is associated with a decrease in the referral stringency of the agency - the evidence yielded by the initial inspection has to be increasingly damning in order to trigger a second round of investigation.

The problem facing the regulator, in formulating an enforcement strategy, is that of choosing an inspection intensity and a referral rule. The choice is, essentially, how far enforcement resources should be directed towards more frequent but less thorough inspections, and how far towards comprehensive but (for that reason) more time-consuming audits. The problem is analogous to that which faces a revenue service designing a tax collection programme. As Roth notes, "(T)he abundance of audit targets creates an ongoing tension between depth versus breadth of audit strategies. The depth strategy argues for thorough and exacting audits. In contrast, the breadth strategy argues for maximising the number of audit contacts by restricting audits to a few simple tests unless a more thorough examination . . . appears warranted" (Roth et al (1988:130). Notice the difference in terminology employed by these authors - in our terminology a 'deep audit' would be what we have labelled an audit whereas a 'shallow audit' would be an inspection. Designating when 'a more thorough examination appears warranted' amounts to fixing a value for Ω. The conflict comes about because of the costliness to the agency of conducting both inspections and audits for ". . . while (regulators) have a broad arsenal of deterrence tools, tight resource constraints demand strategic choices of which enforcement tools to emphasise, in what circumstances, and by what methods. . . [including] depth versus breadth of audit techniques", (Roth et al (1988:129)). It is a recurrent point in this paper that the optimal solution to this strategic choice is likely to be sensitive to the composition of the regulated population - compliers, serious and non-serious violators are different 'beasts' and if uniform policy must be stipulated the relative proportion of each type in the population is likely to be one of the most critical determinants of optimal enforcement strategy.

(3.3) EPA Audit Practices and the 'Polarisation' of Polluters in the U.S.: A Pair of Stylised Facts.

The analysis which we have presented may be helpful in explaining a pair of stylised facts about the recent environmental 'scenery' in the US (the coexistence of which would otherwise be difficult to rationalise), namely that

(i) From the early 1980's onwards the EPA has made a conscious effort to tighten it's referral practices - in our notation to decrease Ω.

(ii) In so far as empirical data and anecdotal evidence are to be believed, during the last ten years patterns of environmental violation have 'polarised'. Many non-serious violators have improved their environmental performance, but this trend has co-existed with a growing trend for 'serious violations' to become a more significant part of the whole.

The policy stance outlined in (i), namely that the EPA's audit effort should be 'beefed-up' (Ω should be reduced) was not accompanied by a parallel effort on the inspection front. Indeed, primary responsibility for random, routine inspection continues to rest with the state-level inspectorates. The EPA constitutes the important 'second tier' of enforcement - it is the threat of audit which leads to those audits being required only comparatively infrequently. William Ruckelhaus (then chief administrator at the EPA) announced in an often-quoted 1983 speech that he wanted the EPA to become the 'gorilla in the closet'; "...the bogeyman that state officials could use as a threat in their dealings with recalcitrant polluters", (Russell (1990:263)). Increasing the potency of that threat amounts to, in our model, a decreased audit trigger level.

Estimating the value of Ω in any given context is extremely difficult. 'Allowable' criminality is not something which the enforcement agency can be expected to announce. Not only could such publication be misconstrued as condoning (rather than merely accepting as inevitable) that degree of criminality, but there are also likely to be enforcement advantages associated with maintaining regulatory uncertainty when the members of the regulated population are risk averse. A 1984 article in the EPA in-house journal *'Inside EPA'* entitled "EPA Will Tolerate 20% to 40% Over Water Permit Levels" put Ω at around one-third in the context of water pollution. In the domain of air quality control the assessment may be 10 to 20%.

The polarisation of firms is catalogued by, amongst others, Russell (1990), who also provides an excellent discussion of the grotesque inadequacies of quantitative compliance records in the United States. He asserts that the frequency and seriousness of comparatively minor violations is decreasing - an encouraging indicator. But the record of serious violations is not: "...when account is taken both of the length and severity of violations as the GAO did when devising a measure of <u>significant non-compliance,</u> the picture is darker" (Russell (1990: 256), underlining not in original, for more detail see Tables 7-4 and 7-5 in Russell (1990: 257-258).

The coexistence of a tightening in audit practices (i.e., a reduction in Ω) and a 'polarisation' of polluting sources is quite compatible with the model in this paper. As we have shown, $(\partial e(NSV)/\partial\Omega)>0$ whilst, in plausible circumstances $(\partial e(SV)/\partial\Omega)<0$. Not only, then, does a decrease in Ω mean that the emissions of a serious violator increase in absolute terms but since $(\partial[e(SV)-e(NSV)]/\partial\Omega)<0$ a reduction in Ω will serve to increase the difference (in terms of emissions) between the representative serious violator and the representative non-serious violator. These comparative static results are more difficult to justify in a model which does not use the 'two-stage' enforcement process which generates the division of violators by type.

(4) Industry Attitudes to 'Inspectability': The Potential Endogeneity of Inspection Noise.

Various authors identify incentives which may exist for the regulated firm to generate 'noise' in the enforcement process. In the context of our model s is the summary measure of the 'noise' in the inspection process, the degree of dispersion of ê around e. In effect s measures how easy or difficult it is for the regulatory inspector to accurately assess the rate of emissions generated by a particular firm, and therefore that firm's compliance status. In plausible circumstances s may be at the discretion, at least to some extent, of the firm - the firm may invest in making itself 'unmonitorable' or untransparent to the regulator. Alternatively it may be able to increase its transparency to the regulator (i.e., generate reductions in its firm-specific value of s) at some cost.

A number of ways in which the firm may be able to increase the 'amount' of noise which the enforcement executive faces in conducting inspections have been identified in the literature. These include choosing technologies which are relatively difficult to inspect, battling court cases or lobbying legislators in order to try to induce rule changes that hinder the cause of meaningful inspection (eg requiring 48 rather than 36 hours' notice to be given by the inspection team, restricting the rights of access of inspectors to different parts of the plant, opposing the constitutionality of the EPA's remote-sensing programme) or obstructing the activities of inspectors in miscellaneous other ways. In the former case, when the burden of proof is on the regulator to demonstrate ('beyond reasonable doubt') that the firm is in non-compliance with any given standard then the firms choice of technology may actually depend upon how 'monitorable' the environmental performance of each competing technology is.

There is much anecdotal evidence that firms in the US have found various ways of reducing the transparency of their activities to environmental regulators. Linder and McBride (1984) label such excercises on the part of the firm as "concealment activity" (Linder and McBride (1984: 327), whilst Kambhu (1989) introduces a variable 'h' which is a measure of "....efforts at deception to hide non-compliance" (Kambhu (1989: 105)). The role of what Harter (1982) terms 'defensive research' (private firms engaging in spiralling amounts of data generation, spurious analysis, excessive documentation with regard to regulatory matters etc.) could be similarly interpreted.

As the examples given suggest, increases in s can be firm-specific or industry-wide. Choice of technique is a firm-level choice whereas the gains from lobbying for restrictions to be put on the inspection rights of the EPA will be industry-wide. For example, the Dow Chemical corporation funded a long-running case against the EPA's use of aerial photography to check for air pollution (by so doing the EPA was able to circumvent the need to give the firm advance warning of inspection, as it would have been obliged to do had it wanted to visit the plant by land). The District and Appeals courts disagreed over whether the use of remote-sensing technologies constituted unreasonable search. Dow forced the issue to the US Supreme Court which decided in a landmark ruling in May 1986 (by a 5 to 4 majority) that aerial surveillance was legitimate, though a distinction was drawn between using standard photographic equipment and more sophisticated sensing instruments. Decisions of this sort constitute important institutional (as opposed to purely technological) determinants of the degree of monitoring noise, s. The changes in s (if any) generated by this type of expenditure would be enjoyed by all firms in the industry. Such expenditure would be most adequately explained, for this reason, in the context of a cooperative model. Free-riding by other firms in the industry upon the expenditures of a company such as Dow Chemical are likely to be endemic to the problem. In the analytic results which follow we concentrate on firm-specific changes in the variable s.

(4.1) When Does a Non-Compliant Firm Dislike Inspection 'Noise' ?

With Propositions (8) and (9) we examine the regulated firm's *incentive* to generate noise. Thus we suppose that there is some way in which the individual firm, by incurring some cost (which may be a direct investment cost, a psychic cost or some valuation of the cost of productive efficiency foregone) can increase the 'noisiness' of its plant. It can 'invest' in increases in s. We ask whether or not it would wish to do so.

In each case we retain the 'small firm assumption', namely that the regulatory parameters are determined exogenously. This is not to say that the incentives for firms to 'invest' in noise are not correlated, but rather that the individual firm is small enough to be able to take regulatory parameters as fixed. In particular the regulator will not adjust Ω or μ (or any other variables under his control) in response to a change in s for a particular firm. We examine two cases; the partial case in which the firm's emissions level is fixed, and the non-partial case in which, having altered s, the firm can adjust its emissions level to its new optimum. The appropriate case in any given context will depend upon the technology in use in the industry, in particular the extent to which the technology permits the plant operator to 'turn-on' or 'off' the flow of effluents or emissions at will. In the terminology of capital theory these could be called the putty-clay and putty-putty cases respectively.

Proposition (8): Ceteris paribus, the non-serious violator prefers decreases in s, the degree of noise associated with the inspection process and the serious violator prefers increases.

Proof: Recall L, the 'loss function' that the firm moves to minimise. It is simple to derive and sign $\partial L/\partial s$ for the NSV and SV respectively;

(21) $L = \{c(e) + \mu.a(e,\Omega,s).\text{Ø}(e,E)\}$

(22) $\partial L(NSV)/\partial s = \mu.a_s(e_{NSV},\Omega,s).\text{Ø}(e_{NSV},E) > 0$

(23) $\partial L(SV)/\partial s = \mu.a_s(e_{SV},\Omega,s).\text{Ø}(e_{SV},E) < 0$

Thus $\partial L/\partial s$ is positive when evaluated for the non-serious violator but negative when evaluated for the serious violator. *QED*

Thus, ceteris paribus, the SV <u>dislikes</u> 'transparency' of the production technology which it uses and would, presumably, have an incentive to hinder the inspection process. The NSV, in contrast, <u>dislikes</u> noise and has an incentive to aid the inspection process and to choose a transparent technology.

The dichotomy follows from the two-stage nature of the monitoring process that we have constructed. By the definition of the non-serious violator $\Omega > e_{NSV} > E$: the NSV is in violation of the standard but not in violation of the audit trigger. Thus if the inspector observes preliminary evidence sufficiently close to e_{NSV} then the firm will not be referred for audit. The NSV therefore benefits from its being 'transparent' to the inspector. The same cannot be said of the SV who is in violation of both the standard and the audit trigger.

The result is rather less determinate when the firm is assumed able to readjust its rate of emissions to take account of the change in regulatory environment associated with the increased 'noisiness' of the inspection process, as the following proposition attests:

Proposition (9): When the firm is able to adust its emissions level, the prefences of both the NSV and SV for increases or decreases in s are ambiguous. In particular they will depend quantitatively and qualitatively upon the form of the inspection uncertainty (that is, the functional form a).

Proof: Completely differentiating the firm's loss function yields

$$(24)\quad dL/ds = \partial L/\partial s + (\partial L/\partial e).(de/ds)$$

$$(25)\quad dL(NSV)/ds = \mu.a_s.\emptyset + \partial L/\partial e.\left\{ -\mu.[a_{es}.\emptyset + a_s.\emptyset_e] / J \right\}$$

$$(26)\ dL(SV)/ds = \mu.a_s.\emptyset + \partial L/\partial e.\left\{-\mu.[a_{es}.\emptyset + a_s.\emptyset_e] \big/ J\right\}$$

In each case the sign of the full-derivative is ambiguous, dependent upon the functional forms employed. *QED*

The result is disappointing for the policy analyst. A priori very little can be said about the incentive facing the representative firm to make investments to increase or decrease their transparency to inspection. The nature of the monitoring uncertainty must be specified to say anything at all. One thing which we can say is that the conventional assertion, namely that a violating firm will always benefit from obstacles and hinderances being put in the way of the regulatory inspectorate, is at best simplistic and at worst qualitatively wrong.

The term in braces on the right-hand side of each expression represents de/ds, the impact of changes in inspection noise, ceteris paribus, upon the optimal choice of emissions by the firm. For both the SV and the NSV it may be positive or negative depending upon the functional form of a and, by extension, of *f*, the qualitative characteristics of the monitoring uncertainty.

This, in itself, is an interesting result. It is contrary to the conjecture of Calfee and Craswell (1984: 999) that reductions in the level of uncertainty in the monitoring process will necessarily improve the the compliance decision of violators. In fact, in a later paper the same authors report simulation results in which reduced uncertainty serves to increase the optimal severity of violation thus contradicting their earlier conjecture (see Table 1 in Calfee and Craswell (1986: 284)). Our results provide analytic support for the ambiguity identified by simulation in their later paper. Linder and McBride (1984) develop a model of monitoring uncertainty in which all non-compliant firms unambiguously gain from an increase in inspection noise. Our results do not support theirs, except as a special case.

(4.2) Optimal Budget-Constrained Enforcement Policy when Monitoring 'Noise' is Endogenous.

In this section we characterise optimal regulatory policy when 's' is firm specific and endogenous. It is demonstrated that for policy to be optimal the choice of the policy mix between audits and inspections must deviate systematically from the standard optimality condition. Changes in the regulatory parameters impact upon emissions behaviour directly, but also induce the regulated population (or some subset thereof) to invest in changes in s (with a consequent indirect change in emissions behaviour). The distortion in the necessary conditions characterising policy optimality captures this class of indirect impacts of regulatory reform. The 'standard' result is seen to be the special case in which s is assumed fixed.

It is assumed that the emissions limit (E) is pre-set (by legislative mandate) and that the planner manipulates the inspection rate (μ) and referral practices (Ω) of the enforcement bureau to minimise aggregate emissions. In employing this objective rather than a more complicated loss function we are applying a dual of the 'Baumol and Oates criterion', which requires that the social cost of satisfying some predetermined environmental quality constraint be minimised. Whilst application of this criterion does not assure Pareto optimality it obviates the need to evaluate increments of damage reduction. In the current paper the enforcement agency minimises aggregate environmental damage subject to some (predetermined) budget constraint, 'ß'. Asterisks are used to denote solution values.

The regulated population is composed of N firms, distributed in some unspecified manner (known to the regulatory agency) according to their incremental compliance costs. The expected expenditure by the agency on applying it's enforcement programme to the 'i'th firm is represented B_i, where $B_i=B_i(e_i,\mu,\Omega,s_i)$. The agency is composed of two divisions - the inspection division and the audit division. Let k_I and k_A be, respectively, the unit costs of conducting an inspection and an audit. Then the total budget of the agency is

$$(27)\ \sum B_i(e_i,\mu,\Omega,s_i) = k_I.\mu.N + k_A.\mu.\sum a(e_i,\Omega,s_i)$$

The first term on the right-hand side is the budget of the inspection division, the second term that of the audit division.

The regulator's problem can, then, be represented

$$(28)\quad \text{Minimise}_{\{\mu,\Omega\}} \sum e_i,\ \text{subject to} \sum B_i(e_i,\mu,\Omega,s_i) \leq ß$$

The various individual rationality constraints (ie the reaction functions of the firm) are, of course, implicit in the analysis but suppressed in the presentation of the problem.

The pair of first-order conditions associated with an interior solution to the regulator's problem are, then,

$$(29)\ 0 = \sum \Big\{ (\partial e_i/\partial\Omega) + (\partial e_i/\partial s_i).(\partial s_i/\partial\Omega) +$$

$$¥.[B_{i\Omega} + B_{is}.\partial s_i/\partial\Omega + B_{ie}.((\partial e_i/\partial\Omega) + (\partial e_i/\partial s_i).(\partial s_i/\partial\Omega)] \Big\}$$

and

$$(30)\ 0 = \sum \Big\{ (\partial e_i/\partial\mu) + (\partial e_i/\partial s_i).(\partial s_i/\partial\mu) +$$

$$¥.[B_{i\mu} + B_{is}.\partial s_i/\partial\mu + B_{ie}.((\partial e_i/\partial\mu) + (\partial e_i/\partial s_i).(\partial s_i/\partial\mu))] \Big\}$$

where ¥ is the Lagrangean multiplier associated with the agency's budget constraint, $\sum B_i \leq ß$.

In the lower equality, the sum of the first two terms represents $de_i/d\mu$, the total impact of a marginal change in inspection intensity upon the emissions of the ith firm. The two components comprise the direct (ie s fixed) effect and the indirect effect. The composite term in square brackets is the total impact of that change upon the agency's budget requirement. The condition dictates that μ should be increased upto the point at which its full marginal impact upon aggregate emissions equals (the shadow value of) its full marginal impact upon agency budget. The upper equality constitutes an analogous condition for the case of marginal changes in the value of the audit trigger, Ω.

Rearranging and eliminating ¥ yields the following condition characterising the optimal 'mix' of the policy instruments, 'inspection' and 'audit'

$$(31) \qquad \frac{\sum[(\partial e_i/\partial\Omega) + (\partial e_i/\partial s_i).(\partial s_i/\partial\Omega)]}{\sum[(\partial e_i/\partial\mu) + (\partial e_i/\partial s_i).(\partial s_i/\partial\mu)]} =$$

$$\frac{\sum[B_{i\Omega} + B_{is}.\partial s_i/\partial\Omega + B_{ie}.((\partial e_i/\partial\Omega) + (\partial e_i/\partial s_i).(\partial s_i/\partial\Omega)]}{\sum[B_{i\mu} + B_{is}.\partial s_i/\partial\mu + B_{ie}.((\partial e_i/\partial\mu) + (\partial e_i/\partial s_i).(\partial s_i/\partial\mu))]}$$

The left-hand side is the ratio of marginal total 'productivities' of Ω and μ (in terms of generating reductions in emissions, the minimisation of which are the policy objective). The right-hand side is the ratio of marginal total impacts upon the binding constraint. The condition is an extension of the simple "ratio of marginal products equals ratio of marginal costs" rule to the case in which the 'transparency' of the regulated population, and thence the marginal productivities of the two policy instruments, are themselves endogenous.

When the 'evasive' behaviour of the regulated population is ignored (that is, the (s x 1) vector s is assumed to be exogenous) it can be seen that the above condition simplifies to the more familiar version, namely

$$(32) \qquad \frac{\sum(\partial e_i/\partial\Omega)}{\sum(\partial e_i/\partial\mu)} = \frac{\sum[B_{i\Omega} + B_{ie}.(\partial e_i/\partial\Omega)]}{\sum[B_{i\mu} + B_{ie}.(\partial e_i/\partial\mu)]}$$

The difference between the two conditions mean that applying the second rule (32) when s is endogenous is likely to lead to sub-optimality. The welfare implications of so doing may be considerable. It is not, however, possible to summarise how optimal enforcement policy, $\{\Omega^*,\mu^*\}$ will deviate from some combination of Ω and μ which satisfy the unadapted condition succinctly. The deviation can be shown to depend qualitatively upon the functional forms describing the monitoring uncertainties, as well as the composition of the regulated population. Furthermore, this dependence is somewhat complex, including second- and third-moment terms. Though this type of conclusion is somewhat unsatisfying from the point of view of the policy analyst, it is characteristic of much of the applied 'second best' literature.

The rationale for the policy deviations is straight-forward to understand, even though the nature of those deviations may be difficult to pin down. Any change in the enforcement régime induces a first round of 'direct' (s fixed) impacts upon industry emissions and agency budget. The firm may also, however, invest in changes in it's 'inspectability' (s) in reaction to the enforcement environment. The implication of this is that there will be a second round of impacts of the policy reform upon emissions and budget. Optimal policy design requires that these indirect impacts be taken account of.

When the firm is assumed to have some discretion over s, the noisiness of the inspection process, it is likely that the expressions derived for the various policy-response derivatives under the assumption that s was predetermined will cease to be valid. They may even change qualitatively. It is important that advocates of regulatory reform take this possibility in to account; firms may react to a supposed 'tightening' of the enforcement régime not by cutting the degree to which they violate standards, but by investing in 'uninspectability'.

(5) Brief Summary and Conclusions.

The aim of this chapter was to investigate some aspects of regulatory reform when there is less than full compliance. In particular, we considered the case of enforcement of an environmental standard (such as an emissions or effluent discharge limit).

A central assumption of our analysis was that of uncertainty in the monitoring process. In particular, it was supposed that the agency is able to conducts two types of checks; inspections and audits. The former are non-exhaustive and yield only preliminary evidence regarding the compliance status of the regulated firm, whereas the latter are exhaustive. The budget-constrained agency uses evidence from randomly conducted inspections to target its audits. An inspection generates a noisy (but unbiassed) estimator of the true emissions of the firm and the agency sets some 'audit trigger' such that an inspection estimate in excess of that trigger causes the inspected firm to be audited.

Exploring the properties of a two-stage enforcement régime of this type was of interest both analytically, and because it constitutes a good stylisation of the way in which a number of environmental protection agencies (notably the United States E.P.A.) do actually function. Our analysis emphasised that, in such a context, it is important to distinguish not just between compliers and violators, but also among violators. The divergence of the audit trigger from the underlying standard means that we can distinguish between serious and non-serious violators. Firms in both groups are in violation of the standards, but only those in the former are so far in violation that an inspection would, in most cases, trigger an audit.

The distinction was shown to be important because serious and non-serious violators will react in different ways to various kinds of policy reform. The efficacy of such reforms can, then, be expected to depend upon the 'anatomy' of the regulated population. This dependence was examined at some length. Of particular interest was the demonstration that non-serious violators will respond to a tightening of the agency's audit referral practices by reducing their emissions, whereas a serious violator may respond by increasing its emissions.

In addition, the two-stage framework was used to reassess the conventional wisdom that a non-compliant firm will always benefit from an increase in the 'noisiness' of the inspection technology which the enforcement agency has at its disposal. We demonstrate that such an assertion is invalid in our model. In general, the serious violator will favour an increase in inspection uncertainty, whereas the non-serious violator will prefer the reverse (though both of these preferences can, in some circumstances, be reversed). The implications of this for the incentive facing different types of firm to invest in 'inspection-proofing' of their operations is obvious.

Appendix.

In the text we defined the distribution $F((\hat{e}-e)/s)$ where s is a strictly positive scaling parameter. The associated density, $f(\hat{e}-e)/s)$ is symmetric and everywhere positive. Furthermore, $f'>0$ when $\hat{e}<e$, $f'<0$ when $\hat{e}>e$.

$a(e,\Omega,s)$ was defined to equal $[1-F((\Omega-e)/s)]$. Of interest to us here are the signs of the various derivatives of the a function in the cases of a Serious Violator (for whom $e>\Omega$) and a Non-Serious Violator (for whom $e<\Omega$). Differentiation of a implies that;

$a_e=f((\Omega-e)/s).(1/s)$, which is positive for both SV and NSV.

$a_\Omega=-f((\Omega-e)/s).(1/s)$, which is negative for both SV and NSV.

$a_{e\Omega}=f'((\Omega-e)/s).(1/s^2)$, which is negative (positive) for a NSV (SV).

$a_{ee}=-f'((\Omega-e)/s).(1/s^2)$, which is positive (negative) for a NSV (SV).

$a_{\Omega\Omega}=-f'((\Omega-e)/s).(1/s^2)$, which is positive (negative) for a NSV (SV).

$a_{es}=-f'((\Omega-e)/s).((\Omega-e)/s^2).(1/s)-f((\Omega-e)/s).(1/s^2)$, the sign of which is ambiguous for both NSV and SV.

$a_s= f((\Omega-e)/s).((\Omega-e)/s^2)$, which is positive (negative) for a NSV (SV).

Environmental Standard-Setting With Endogenous Technical Change: A Dynamic Model and Application to the Nuclear Electricity Sector.

(0) Introduction.

The principle that environmental risk from nuclear power plants should be kept 'As Low As Reasonably Achievable' (ALARA) is deep-rooted in the regulatory principles and practices of a number of countries (including the US, Canada and Spain). In its simplest formulation the ALARA rule dictates that at any moment the probability of an accident at a nuclear plant should be set at that level at which the social marginal costs and benefits of additional increments of risk reduction are exactly balanced. It is an embodiement of the classical cost-benefit decision criterion familiar from economics text-books and, for the reasons given therein, is an attractive policy rule.

ALARA is a particular case of technology-based standards - the stringency of safety regulations at any particular time is determined by the current state of technology. Whilst this flexibility is a supposed strength, it is also its critical weakness. The 'moving goal-post' regulations which it generates in a dynamic environment will, we argue, inhibit the supply of new externality-control techniques. These inter-temporal inefficiencies are may swamp the intra-temporal efficiency gains which the flexibility generates.

In this paper we seek to examine the optimality of ALARA regulation in a dynamic context in which the potential for technological advance is great. Our hypothesis is straight-forward. Whilst the ALARA rule has impeccable optimality credentials as a basis for policy in a static regulatory environment, these credentials will generally not be applicable in a dynamic setting.

The ALARA rule dictates that regulatory requirements should be ratchetted upwards as new safety-technologies become available. In an industrial environment in which technical advance is endogenous (ie the regulated industry is expected to invent and commercialise new safety designs) the anticipation of such ratchetting can be expected to adversely affect the incentive that industry has to innovate, and thus to diminish the supply of new techniques. Dynamic optimality of a policy requires that the policy-maker take these dynamic considerations into account - the innovations of today provide the technical basis for the increased stringency of standards tomorrow. The regulator's problem is to decide when and by how much to tighten standards, taking into account the detrimental effect which such tightening can be expected to have upon the rate of technological advance.

In Section 1 we will consider the role of the ALARA rule, and variants of it, in the determination of safety regulation in the nuclear sector internationally. We will also emphasise its key characteristics. Section 2 surveys some of the contibutions which investigate the impact of environmental regulation upon the supply of new externality-abatement technologies. Though we apply the model to the nuclear sector the results are clearly generalisable to any regulated industry which generates a negative externality. Thus 'emissions' can be interpreted in the broadest sense of the word. A nuclear reactor, for example, 'emits' accident risk to the surrounding countryside and population just as surely as a coal-fired plant emits smoke.

In Section 3 a simple static model of environmental standard-setting in the nuclear industry is set-up. The optimal policy rule is derived and noted to be the ALARA rule.

In Sections 4 and 5 a simple two-period model is used to examine the regulated industry's incentives to devise and commercialise new safety technologies under different regulatory regimes. Thus the dynamic behaviour of the industry is characterised. Optimal standard-setting is re-examined subject to this dynamic incentive compatibility constraint. Sections 4 and 5 differ by the assumed source of technological advance. In Section 4 improvements in safety technology come about as a result of the conscious R&D efforts of the industry. In Section 5 improvements are the result of 'learning-by-doing' - licensees 'learn' to be safe through accumulating operating experience. In a sector employing technology as sophisticated as that used in the nuclear energy sector, much is learned about how to exploit and contain the technology from near accidents (such as those at Three-Mile Island, Windscale and Browns Ferry) and other smaller stochastic events. Such learning comes about if operational experience increases. In Section 6 we consider the interaction of learning by doing and R&D driven technological improvements. We introduce the concept of 'Experience-Directed Research and Development' (EDR&D) and argue that it is a good stylisation of technological advance in the nuclear power industry.

In each case the behavior of the representative firm is described and optimal regulatory policy characterised. We conclude that in this type of dynamic setting the regulator will generally be able to increase social welfare by departing from the ALARA formula. The nature of the departure will, however, depend <u>qualitatively</u> upon whether learning is primarily of the 'by doing' type or results from distinct research and development decisions. It will also depend upon the extent to which the regulatory bureau is able to commit itself to future regulations.

(1) The 'ALARA' Rule in National and Supra-National Regulation of the Nuclear Industry.

Policy documents and academic papers too numerous to mention exult the normative superiority of the 'As Low As Reasonably Achievable' (ALARA) rule in the field of environmental standard-setting. The underlying concept of ALARA is that the regulator should estimate the marginal cost and marginal benefit schedules associated with successive increments of risk-reduction, and determine risk requirements by considering the juxtaposition of the two[1]. To convert this consideration into a policy 'rule' requires stipulation of a weighting factor - what relationship does the planner wish to exist between marginal cost and marginal benefit? Textbook, Benthamite cost-benefit analysis would assert that the appropriate weighting factor be unity. In this case the point of interest for the planner would be the point of intersection between the two marginal curves. In most cases where the ALARA methodology has been applied, however, it has been combined with a factor which advocated disproportion between costs and benefits, with the disproportion weighing in favour of 'excessive' (vis-à-vis the Benthamite ideal) risk-reduction.

The ALARA rule, and variants on it, have received considerable acceptance in the world of nuclear industry environmental regulation.

The main tenet of the International Commission on Radiological Protections' (I.C.R.P) guidelines on dose limitations is that "...all exposures shall be kept as low as reasonably acheivable (ALARA), economic and social factors being taken into account" (ICRP (1977)). In fact, these guidelines constitute the revised version

(1) With particular reference to the nuclear industry, a good cross-section of these can be found in the proceedings of a symposium on "ALARA - Quantitative Optimisation Techniques for Radiation Protection in the Nuclear Industry" held in London, England in 1986, published in their entirety (with supplementary commentary) in "ALARA. Principles, Practice and Consequences" (Lakey and Lewins (1987)).

of those originally laid down by the ICRP in 1965 that doses be kept 'As Low As *Readily* Achievable' (italics not in original) (ICRP, 1965, Paragraph 52). These were updated in 1973 on the basis that (according to Webb (1987)) "...the Commission concluded that the second word more closely described its intentions than the first word" (Webb, (1987: 4). The only way in which we can interpret this is that some suggested weighting factor (of costs and benefits at the margin) is implicit in each word, and that the ICRP preferred the factor implicit in the word 'reasonably'. The fact that the weighting factor to be adopted is not specifically enunciated renders the ICRP guideline vague at best.

Whilst formal jurisdiction over nuclear safety matters resides with the national member governments, the European Community also invokes the ALARA criterion in setting forth its position on nuclear sector safety and environmental issues. The Council of Ministers current Directive dates from July 1980 (subject to a few non-substantive amendments adopted in 1984) and stipulates that "....(T)he limitation of individual and collective doses resulting from controllable exposures shall be.....kept A.L.A.R.A." (See Official Journal L264/1 [80/836/Euratom]). In the related domains of potentially harmful chemical, physical and biological agents (including non-radiological carcinogens) the analogous Directive (80/1107/EEC) is that exposures be kept 'As Low As Is Reasonably Practicable' (A.L.A.R.P.) providing that they are not also unreasonable. The nomenclatural distinctions routinely employed in these documents are, we hope, more meaningful to the lawyer than they are to the economist (see Kaser (1989)).

As well as its hallowed place in international standard-setting ALARA (and its hybrids) has also taken a firm grip on the national regulatory philosophies of a number of countries. A noteworthy example is the United Kingdom. The United Kingdom regards itself as the pioneer of the ALARA methodology and, it must be said, is rather proud of that heritage. As the Chairman of the Central Electricity Generating Board (CEGB), Dr J. K. Wright, asserted in his preamble to the 1986 conference cited earlier "......I want to set the tone which I know our speakers will express in more detail, that ALARA is essentially a British concept....(it is) *our* approach to regulation", (Lakey and Lewins (1987: 2))[2]. The precedential interpretation of the ALARA criterion is that a risk is satisfactory if there is a 'gross disproportion' between the marginal cost and benefit of additional risk reduction. Inspection of the evidence presented to the Sizewell planning inquiry would suggest that, at least as a rule of thumb, a factor of 10 to 1 is interpreted as satisfying this requisite.

---

(2) For a detailed survey of the role of the ALARA criterion (and the various hybrid versions of it) in the nuclear electricity sector in the United Kingdom (and more generally) see Webb (1987). The legal principle underlying ALARA predates, however, its application to the regulation of nuclear hazards. The principle that industry should use 'best practicable means' (commonly interpreted as being synonomous with ALARA) was enumerated in the UK air pollution regulations of the nineteenth century (the Alkaline Acts, 1863). It became defined by statute in the UK Clean Air Act (1956, Section (34.1)) and the UK Control of Pollution Act (1974). Space precludes a detailed history of judicial interpretation of ALARA and related considerations. The interested reader is directed to the excellent legal analysis, covering several European countries (including the UK), provided by Kaser (1989): "Acceptable Nuclear Risks: Some Examples From Europe".

The central characteristic of the ALARA principle (of which ALARP is one common variation) is that current safety standards are determined by inspection of current technologies. ALARA regulators are frequently said to employ 'optimising methods'. The problem examined in this paper is that those optimisations are performed subject to the existing technological constraints. Such static optimisations are most unlikely, we demonstrate, to provide adequate incentives for the industry to loosen those technological constraints. Applying myopic ALARA rules through time reduces the essentially dynamic problem of regulating an industry employing an immature technology to a sequence of momentary, static problems. We examine optimal solutions to the *dynamic* problem explicitly, under various assumptions regarding the underlying regulatory environment.

(2) Environmental Standard-Setting and Innovation: A Selective Survey.

A number of authors have attempted to assess how environmental regulation may impact upon the innovation and diffusion processes at work in the regulated sectors and outside. Of particular focus have been the comparative performances, in this regard, of different regulatory styles; emissions standards, emissions subsidies and/or charges, and tradeable emissions permits.

The consensus of the results is that direct controls (of the type to be examined in this paper) generate less innovation in abatement technologies than do the other instruments, at least in deterministic models.

Wenders (1975), Zerbe (1970) and Orr (1976) demonstrate that emissions charges induce more innovation than do direct controls. Similarly, Downing and White (1986) use a diagrammatic approach to rank the dynamic performance of four alternative regulatory instruments. They find that emissions charges can be expected to generate the most innovation and emissions standards to generate the least. Milliman and Prince (1989) develop broadly similar results (though qualified for the case in which the innovating firm is able to extract royalties form other firms in its industries which install the innovation), as do Marin (1978), Kneese and Schulz (1975), Russell (1979) and Mills and White (1978).

Other authors have identified ambiguities in the dynamic rankings of the alternative régimes, however. Wesley Magat, in a pair of papers (Magat (1978, 1979)), constructs a dynamic model of abatement-targetted R&D by the firm and identifies cases in which direct controls may generate *more* research effort than any of the other policy instruments considered. Mendlesohn (1986) presents a model incorporating research uncertainty, concluding that quantity-based rules will generally encourage more efficient levels of technical change than will price-based rules. In so doing Mendlesohn (1986) extends the classic result of Weitzman (1974) to the case in which technical change is endogenous.

Of critical importance in assessing the innovation and diffusion incentives in a particular context is the assumption made about how the regulatory agent is expected to respond to technical change. The ALARA rule, for example, characterises one algorithm according to which emissions standards could be updated in response to changes in the technological environment. The contributions of Zerbe (1970), McHugh (1985) and Magat (1979), amongst others, assume away agency response. Magat assumes that "...no single firm can influence either the type of environmental regulation imposed upon it or the levels of the regulatory parameters". (Magat (1970: 7)). Such a comprehensive 'small firm' assumption effectivly sweeps away the dynamic-incentive problems which are the focus of the present paper and in so doing may obscure much that is of interest in the discussion of optimal environmental protection. Yao (1988) provides an excellent model of regulation in the automobile industry in which an oligopilised industry and a regulatory agency react and respond to each other in a two-period game. It is this type of model which we present in the current paper. Unlike Yao, we examine optimal policy under three different assumptions regarding the ability of the agency to commit to future standards, and we also consider the possibility of learning-by-doing.

The focus of the present paper is not, it should be noted, to compare the innovative performance of the regulated industry when the regulator employs alternative policy instruments (standards, taxes or whichever). It is, rather, to examine the optimal employment of a particular instrument, namely the emissions standard. We assume that the regulator is constrained to set emissions standards and characterise how he should do so in order to optimise a well-defined social goal function. Within these constraints we question the efficacy of the particular criterion, namely the ALARA rule, which has become increasingly popular as a solution to the particular problem which we address. The question is a reasonably simple one in a static context, with available technology either fixed or developing according to some exogeneously determined pattern. It is less straightforward when the supply of abatement technologies is endogenous to the regulatory process.

The 'conventional wisdom' in most of the environmental policy literature of the last few years is that the ALARA rule (which boils down, essentially, to application of a static cost-benefit rule through time) is an appropriate basis for determination of standards. It is this conjecture which we examine in the model which follows. In particular, we emphasise the non-robustness of the ALARA rule's efficiency characteristic to a dynamic context

(3) Static Model of Safety Regulations.

The static efficiency of the ALARA principle is trivial to demonstrate in a single-period setting.

Producer's surplus or profit is denoted by P. The public and environmental risk associated with the industry is denoted D. The regulator is risk neutral and the aim of regulation is to maximise a weighted sum of these; SWF=μ.P+D, where μ is the relative weight that the planner puts upon producer's surplus.

Consider s, the complement of the probability of a plant suffering an accident in a given period. The variable s can loosely be regarded as the 'amount of safety', scaled on the positive unit interval, $0 \leq s \leq 1$. When s=1 the plant is absolutely safe. The gross benefit to the producer from operating a plant for a period is assumed to be a constant, b. It is assumed that the cost of operating a plant for a period, c, depends upon the accumulated amount of R&D, r, and the safety level at which that plant is operated. Other arguments of the cost function are surpressed. Thus c=c(r,s). It is assumed that the function is at least twice differentiable in both of its arguments. In this section the static nature of the problem means that we treat the value of r as fixed and exogeneous. Letting subscripts denote derivatives in the usual way, it is plausible to suppose that the cost function is concave in r and convex in s: $c_r<0$, $c_{rr}>0$, $c_s>0$, $c_{ss}>0$. Diminishing marginal returns to R&D expenditures imply that c is concave in r. Increasing r lowers the cost of achieving a given level of s. In addition, we assume that $c_{rs}<0$ - the higher is the level or r the lower is the <u>marginal</u> cost of achieving a particular improvement in safety. Equivalently, the marginal benefit of research decreases with the level of regulatory stringency. Costs will be increasing in safety levels, with an asymptote at s=1 (Figure 1).

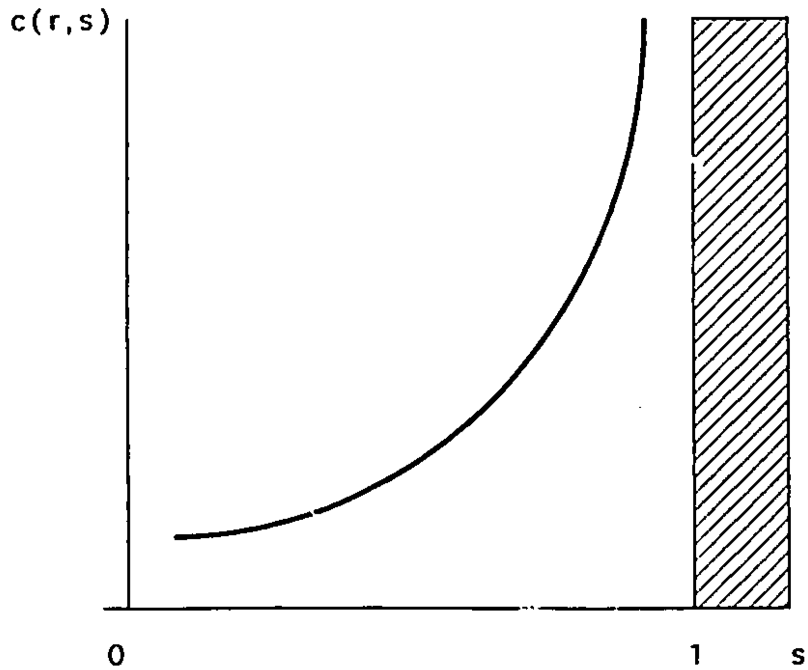

Figure 1: The assumed relationship between plant safety (s) and per period operating costs, c(r,s).

When a plant suffers an accident the damage done is d per unit of capacity (e.g. per kW). Note that we consider, in the term d, only that portion of accident costs not incident upon the firm, ie that portion of off-site costs which constitute the 'externality' associated with an accident. It is this externality which constitutes the source of the market failure which, in turn, provides the rationale for regulatory intervention.

The regulator's problem is to set a safety standard (which is, by assumption, fully and costlessly enforced), s, to maximise single-period social welfare. It is assumed that the optimal value of s, s*, will lie strictly between 0 and 1. Clearly, then D=d.(1-s). This problem and the associated first- and second-order conditions (FOC and SOC) which characterise the interior solution are, respectively;

(1) $\text{Maximise}_s \left\{ \mu.[b - c(r,s)] - d(1-s) \right\}$

(2) $s^*: \partial SWF/\partial s = -\mu.c_s(r,s) + d = 0$

(3) $s^*: \partial^2 SWF/\partial s^2 = -\mu.c_{ss}(r,s) < 0$

The FOC (Equation 2) dictates that the regulator requires risk reduction at the representative plant up to the point at which the social valuation of the cost of further increments exceeds the social gain. This is merely a re-statement of the ALARA rule oulined in the introduction to this paper. The second-order condition will always be met, given our assumption that costs are convex in s. Notice the obvious comparative statics; the more damaging is the expected accident or the lower the weight placed on producer surplus, the tighter will be regulatory standards. In particular;

(4) $\partial s^*/\partial d = 1/\mu.c_{ss}(r,s) > 0$

(5) $\partial s^*/\partial \mu = -c_s(r,s)/\mu.c_{ss}(r,s) < 0$

In this section we have confirmed the familiar (almost tautological) optimality of the ALARA rule in standard-setting in a static regulatory environment. The optimality of the rule will be sustained in the trivial pseudo-dynamic context in which the supply of new techniques is exogeneous. In such a case the periods are analytically separable and the planner can apply the ALARA criterion to set standards on a period by period basis. The 'dynamic' problem in such a case is, in effect, just a sequence of discrete static problems.

The problem with drawing policy conclusions from such analysis is that it fails to take account of the intertemporal aspects of standard-setting when the supply of new techniques in endogeneous. In this case myopic application of the ALARA algorithm can yield unsatisfactory policy results.

(4) Simple 2-Period Model of Safety Regulations With R.&D.

As in the last section, the costs of operation for the firm in a particular period depend upon the stringency of regulations that must be met and upon r, the accumulated volume of R&D expenditure.

Throughout the paper we will refer to the 'industry' and the 'firm' inter-changeably. One simple justification would be to assume a monopolised electricity supply industry. This assumption is not, however, necessary. It is straightforward to demonstrate that the aggregative results are not qualitatively affected by market structure. They are independent of whether we envisage an industry monopolised by one firm, comprised of n symetric firms each constituting (1/n)th of the industry, or comprised of some combination of heterogeneous firms. The direction of the aggregative Nash equilibrium response by a non-monopolised industry to regulatory parameter changes can be shown to be 'as if' the industry were indeed monopolised.

In a two-period model the firm's problem is to choose a level of R&D expenditure in the first period. The first-period R&D induces a downward shift in 'cost of safety' curve faced by the firm in the second period.

(4.1) The Firm's R&D Budgeting Decision.

The regulator is permitted to update regulations at the start of each period. The firm will therefore face a sequence of safety requirements, s_1 and s_2, where the subscript denotes the period in which the standard applies. The firm's problem is to choose an optimal level of R&D expenditure for the first period, r^*, given the regulatory environment $\{s_1,s_2\}$ and a discount rate §.

$$(6)\quad \text{Maximise}_r\, P{:}\left\{ b.[1+ \S] - [c(0,s_1) + \S.c(r,s_2)] - r \right\}.$$

For simplicity we will assume throughout that §=1 (ie the no discounting case). This improves the clarity of the paper without significantly changing the principal results. The first- and second-order conditions asssociated with its choice of r^* are therefore;

$$(7)\quad r^*{:}\ -\partial P/\partial r = 1 + [c_s(r^*,s_2).(\partial s_2/\partial r) + c_r(r^*,s_2)] \leq 0$$

$$(8)\quad r^*{:}\ \partial^2 P(r^*)/\partial r^2 = -\ 2.c_{rs}(r^*,s_2).(\partial s_2/\partial r) - c_s(r^*,s_2).(\partial^2 s_2/\partial r^2) - c_{rr}(r^*,s)$$

The second-derivative will be denoted by J for brevity, where we assume that the parameter values are such that J<0 to ensure that the turning point identified is a maximum. The first-order condition dictates that the safety R&D budget, r, will be increased up to the point at which the value of the marginal dollar is equated with the present value of the incremental saving to the firm of gross period 2 compliance costs. The term in square brackets in the foc is the full derivative of period 2 (unit) costs with respect to a change in r - composed of the impact through the regulatory changes that the research effort induces, and the s_2-constant effect respectively. If this square bracket term is positive (if increases in r serve to increase gross second period compliance costs) this is a sufficient (but not necessary) condition for the corner solution, $r^*=0$, to dominate. In this case industry will prefer not to do any R&D and the pool of safety technology available will stagnate. If the regulator is permitted to update regulations at the start of the second period then $\partial s_2/\partial r$ may be non-zero. It is plausible to suppose that $\partial s_2/\partial r \geq 0$ - more R&D on safety technology will not reduce the stringency of safety standards in future periods. Thus, in spending money on improving safety technologies the firm can expect to induce stricter second-period regulation.

Assuming that the planner is unable to coerce the industry to engage in safety R&D (the monitoring problems associated with enforcing such a requirement could be expected to be prohibitive), the firm's R&D decision-making calculus constitutes an incentive constraint upon the planner's problem. Cost-justified reductions in environmental risk in period 2 can only be instigated insofar as technology allows. This, in turn, requires that industry engage in period 1 R&D expenditures. The propensity of the industry to engage in research activities depends, however, upon how the regulator will 'use' the results of that activity. If the regulator responds to a high level of R&D by large increases in safety requirements in later periods (ie regulations are highly 'reactive' in this sense) then this reduces the incentive that industry has to engage in that activity. This can be demonstrated more concretely, momentarily treating $(\partial s_2/\partial r)$, by noting that the following derivatives are implicitly defined;

(9) $\partial r^*/\partial(\partial s_2/\partial r) = c_s(r,s_2)/J < 0$

The planner's problem is to choose how to set and update standards, given that he is subject to the above incentive constraint. The cost of R&D effort is borne entirely by the industry and generates a surplus. The surplus is left for the firm if regulations are not tightened; in that case all of the surplus is taken as increased producer's surplus - the industry enjoys lower period 2 costs in attaining unaltered safety standards. Alternatively, period 2 safety requirements may be tightened such that some, all or more than all of this potential increase in producer's surplus is converted into improved environmental quality in period 2. The greater is $(\partial s_2/\partial r)$ the greater the extent to which this conversion occurs.

(4.2) The Supply of Research Effort Under an ALARA Régime.

The ALARA rule is one convention which the regulator can choose to use in updating regulations. That is to say it is one (of many) conventions which can be adopted for the division of the surplus mentioned above. It is our contention that, when the impact of the division rule upon the supply of new technologies is taken into account, the ALARA rule is unlikely to be the optimal policy régime.

In an ALARA world regulatory standards are continually updated so as to equate the incremental (static) costs and benefits of further impovements in safety. Under such a régime, therefore, standards in any period are implicitly defined as a function of cumulative R&D investment - $s_i=ß(r_{i-1})$, $ß'>0$, $ß''<0$. In the 2-period case, $s_1=ß(0)$, $s_2=ß(r)$, satisfying;

$$(10) \quad s_1: c_s(0,s_1) = d/\mu$$

$$(11) \quad s_2: c_s(r,s_2) = d/\mu$$

ALARA means that standards are continually updated to ensure that c_s is kept equal to d/μ. Since r induces downward movements in the 'marginal cost of safety' curve (see Figure 2), regulations are tightened as shown. Period-2 requirements are mechanically related to period-1 R&D budgets. The regulator 'creams off' some portion of the surplus generated by the firm's research efforts. Note that, if area(abcd)>area(aefg) then period-1 R&D actually increases the period-2 costs of the firm. In that case the

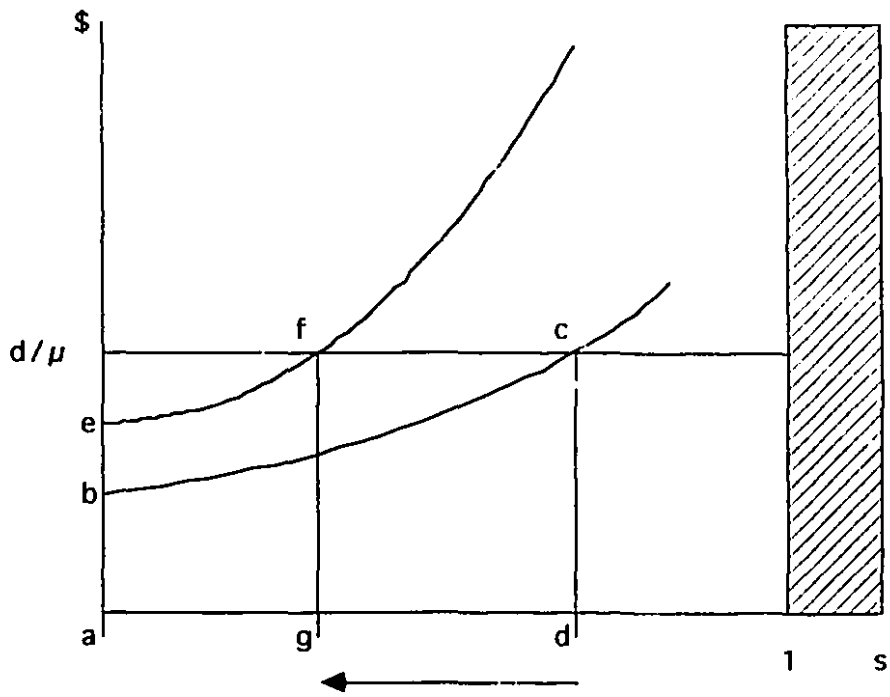

Figure 2: Downward shift in marginal cost of safety curve generated by R&D. ALARA criterion dictates that regulatory requirements be tightened from d to g.

regulator is so aggressive in updating standards in response to technological advance that the cost associated with meeting the additional requirements is greater than the cost saving of being able to satisfy the previously existing regulations more cheaply. In this case the corner solution, $r^*=0$, is the firm's reaction.

The first- and second-order conditions associated with the firm's choice of R&D budget can now be written more compactly for the case where it faces an ALARA regulatory regime;

$$(12)\quad r^*: -\partial P/\partial r = 1 + c_s(r^*,s_2).\beta'(r^*) + c_r(r^*,s_2) = 0$$

$$(13)\quad r^*: \partial^2 P(r^*)/\partial r^2 = -2.c_{sr}(r^*,s_2).\beta'(r^*) - c_s(r^*,s_2).\beta''(r^*) - c_{rr}(r^*,s) < 0$$

It is straightforward to demonstrate that the firm will not do 'enough' R&D by evaluating the sign of $\partial SWF/\partial r$ at r^*;

$$\begin{aligned}(14)\quad \partial SWF/\partial r &= -\mu.[\partial P/\partial r(r^*)] + d.\beta'(r^*)\\ &= b.\beta'(r^*)\\ &> 0\end{aligned}$$

The term in square brackets is zero by application of the envelope theorem to the firm's problem. Figure 3 serves to clarify the situation.

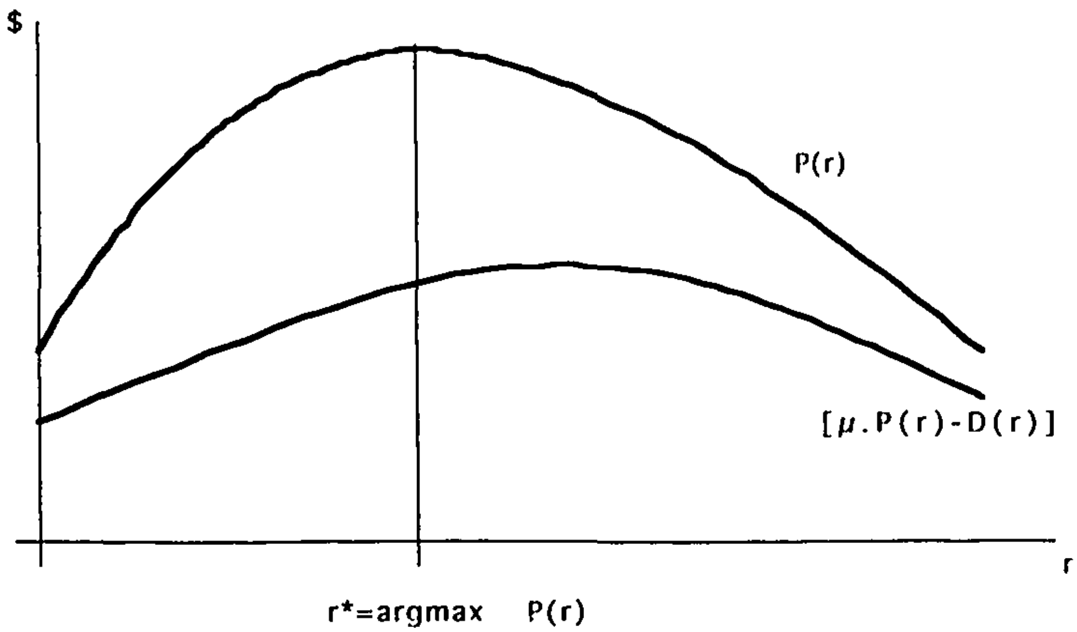

Figure 3: Industry surplus and social welfare as functions of amount spent on safety-related R&D. Socially optimal level exceeds that whic is privately optimal under ALARA regulation.

Thus we have a classic public-goods problem; under ALARA regulation a proportion (and that proportion may be more than 1) of the returns to R&D in safety techniques is accrued outside of the industry. The incentive facing the industry to engage in this socially beneficial activity is, therefore, lower than optimal[3].

Assuming that the regulator's only policy instrument is his authority to set safety standards (he cannot, for instance, subsidise safety research directly) then we have an *a priori* justification for adjusting regulation in the face of dynamic incentives. In the following sections we attempt to characterise the 'optimal' departure from ALARA in this context.

(4.3) Optimal Standard-Setting by Stackelberg Leading Regulator With Commitment.

In this section we consider the case where s_2 is fixed when r varies. The regulator is assumed to commit to period 2 regulations at the start of the game. His problem is to determine what level of stringency he should commit to. The equilibrium sought is thus the open-loop solution to the game in which temporal rationality of the planner's strategy is not required. Analytically this implies that $\partial s_2/\partial r=0$. The firm's R&D problem and associated first-order condition can then be represented;

---

(3) Cabral and Riordan (1989) and Clemhenz (1991) identify an analogous problem in the context of 'price-cap' regulation. If the cap is periodically revised to reflect cost-saving advances in the production technology used by the industry, then the industry's incentive to strive for such technological improvements will be diminished.

$$(15)\quad \text{Maximise}_r\, P: \left\{ 2b - [c(0,s_1) + c(r,s_2)] - r \right\},$$

and

$$(16)\quad r^*: -\partial P/\partial r = c_r(r^*,s_2) = -1$$

Equation 16 implicitly defines the level of R&D expenditure chosen by the firm as a function of the stringency of the second period standards to which the regulator commits himself. In particular,

$$(17)\quad \partial r^*/\partial s_1 = 0$$

$$(18)\quad \partial r^*/\partial s_2 = -c_{rs}/c_{rr} > 0$$

The firm's choice of a research budget is invariant to changes in the stringency of first period regulations and increasing in the stringency of those of the second period.

The regulator's problem is to choose a pair of requirements, $\{s_1,s_2\}$, to maximise social welfare, subject to these incentive compatibility constraints. As in the earlier case the SWF is the social valuation of the stream of producer surplus minus the stream of environmental damage associated with the operations of the industry;

(19) Maximise SWF: $\{\mu.P(r^*) - D\}$
$\{s_1,s_2\}$

or, more explicitly, to choose a sequence of standards $\{s_1,s_2\}$ to

(20) Maximise $\Big\{\mu.[2b - [c(0,s_1) + c(r,s_2)] - r]$

$- d.[2 - s_1 - s_2]\Big\}$

subject to $[1 + c_r(r^*,s_2)] = 0$

The two necessary conditions characterising the interior solution to the the policy problem are, then;

(21) $\partial SWF/\partial s_1 = -\mu.c_s(0,s_1) + d = 0$

(22) $\partial SWF/\partial s_2 = -\mu.[c_s(r^*,s_2) + c_r(r^*,s_2).(\partial r^*/\partial s_2)]$
$-\mu - d = 0$

These optimality conditions can be rearranged with c_s on the left-hand side, to make them easy to compare with the ALARA regulatory rule which dictates that $c_s=(d/\mu)$ in each period:

(23) $c_s(0,s_1) = d/\mu$

(24) $c_s(r^*,s_2) = d/\mu - [1 + c_r(r^*,s_2).(\partial r^*/\partial s_2)]$

$$= d/\mu - [1 - c_r.c_{rs}/c_{rr}]$$

These two equations determine optimal policy and their interpretation is of considerable interest. We can examine the policy implications of our analysis for the two periods separately.

The first equality demonstrates that the regulator should apply the ALARA formula in period 1. Optimal policy will not involve regulatory relief to industry whilst safety technologies remain immature.

The second equality shows that, in the general case, period 2 standards will not be updated to those levels which would be dictated by the ALARA algorithm. They will do so only in the particular case in which $c_r(r^*,s_2).(\partial r^*/\partial s_2)=-1$. Convexity of costs in s ($c_s,c_{ss}>0$) implies that, in period 2, industry should be offered regulatory relief (vis-à-vis ALARA) iff....

$$(25)\quad [1 + \underset{(-)}{c_r(r^*,s_2)}.\underset{(+)}{(\partial r^*/\partial s_2)}] > 0,$$

if, that is, $(\partial r^*/\partial s_2)$ is sufficiently large in the neighbourhood of equilibrium. The condition can, after substitution, be rewritten

(26) $[c_r . c_{rs}/c_{rr}] < 1$

If the inequality in the condition is reversed then period-2 regulations should be <u>tougher</u> than their ALARA counterparts.

Optimal regulations in the Period 2 will be stricter (laxer) than dictated by ALARA when $[c_r . c_{rs}/c_{rr}]<1$ (>1). The ambiguity is to be expected. As we argued earlier, ALARA regulation will encourage an inefficiently low flow of R&D funds. With the ability to commit, the regulator can rectify this by commiting to non-ALARA Period 2 standards in order to encourage extra innovation. The above condition merely dictates whether the correct way for the regulator to accomplish this is through stricter period-2 standards or laxer period-2 standards[4].

Given that $(\partial r^*/\partial s_2)>0$, it is possible but not certain that tightening standards is the appropriate policy rule. The belief that this is the case is the rationale behind the 'technology-forcing' regulatory statutes in US mobile source (i.e. car) emissions regulations[5]. The regulatory treatment of this sector in the

---

(4) With $\partial c/\partial r$ and $\partial(\partial c/\partial r)/\partial r$ fixed at some arbitrary values, a sufficient condition for the former to be appropriate is that $(\partial^2 c/\partial r . \partial s)$ be 'sufficiently' large, that is that the impact of an increased research budget upon the marginal increase in compliance cost associated with tighter standards be sufficiently large.

(5) It should be noted that there is a distinction between 'technology-forcing' of the type modelled in this paper and that envisaged by, amongst others, Ashford and Heaton (1983).

1970's provides a classic example of this and some related issues. The 1970 amendments to the US Clean Air Act mandated hydrocarbon and carbon monoxide emissions reductions on all new motor cars of 90 % by 1975, with a similar reduction in nitrogen oxide required by 1976[6].

To recapitulate on the analytic results of this section: Optimal policy will <u>not</u> involve regulatory relief being granted to the industry during the Period 1. Regulations in the Period 2 will be laxer/stricter than dictated by ALARA when $c_r . c_{rs} / c_{rr}$ is less than/greater than unity respectively. Stricter regulations will be selected (ie 'technology-forcing employed') when an increase in environmental regulations have a sufficient impact upon the marginal productivity of the incremental research dollar.

(4.4) Optimal Standard-Setting by a Stackelberg Leading Regulator Without Commitment.

It is easy to see and to show that without commitment the regulator can do no better than imposing myopic ALARA

In the current paper, as the reader will have noted, a technological advance means a change in the marginal cost conditions associated with externality reduction. The range of *feasibility* of the technology is not extended. Throughout this paper any technology will be associated with a marginal cost function defined for all values of s in the half-open interval [0,1).

(6) In the event, the requirements and schedules for compliance were revised in favour of the automobile industry in 1973 and 1977 when, having arrived at the critical hour without having made the necessary technical advances, industry chiefs threatened to shut-down assembly lines rather than produce illegal cars and risk fines of $10,000 per vehicle. The problem with this, and various other, attempts to 'force' technological advance in the domain of environmental protection has been the difficulty that the relevant institutions (including Congress) have had in committing *credibly* to future standards.

regulations in both periods. The result follows easily by backwards-induction. Having arrived at the start of the mature phase (period 2) the range of technologies developed is predetermined and the regulator can do no better than use the ALARA criterion. If he wants to offer regulatory relief, therefore, the only way in which he can do it credibly is in the learning phase itself (period 1). We noted in the last section, however, that the firm's R&D decisions do not depend on current but on future regulatory standards (in particular, $\partial r^*/\partial s_1=0$). Period one regulatory relief therefore has no incentive its R&D path, and so the regulator can do no better effect, laxer period one regulations do not induce the firm to alter than employ the ALARA rule in this period also.

Thus, without commitment, the outcome is the same as it would be if the regulator were appointed Stackelberg follower (and thus obliged to decide upon regulatory stringency <u>after</u> the industry had committed itself to a programme of research for the first period).

Without the ability to commit the regulator is stuck with the inefficient two-period ALARA outcome. We would anticipate insufficient first-period R&D effort from the industry and thus a mature phase characterised by the use of 'under-'developed (compared to the social optimum) technology. To improve upon this situation the regulator must find some alternative policy instrument; susidy of R&D, coercion of firms into innovative effort, and the direct control of the R&D programme at 'government laboratories' are three possibilities.

(4.5) Optimal Standard-Setting with Partial Commitment: Optimal Grandfathering.

We have seen that the ability of the regulator to effect outcomes socially superior to that associated with myopically following ALARA rules in each period depends very much upon his ability to commit to what he will do in the future periods. In effect he wants to be able to say to the industry; ". . . do your utmost to develop safer and cheaper technologies and I promise that I'll leave you a sufficient portion of the social gains from those technological advances to make your efforts worthwhile".

It is reasonable to assume that, typically, the regulator's ability to commit is not well described by either of the polar cases considered above. In this section we look at the case in which 'grandfather clauses' are feasible. The 'raw' regulations which must be applied to any new plant at time t are determined by the ALARA algorithm. The regulator has discretion, however, over the extent to which existing plants are required to be updated to meet the newer, presumably more stringent, safety standards. Thus we use the term 'partial' commitment to refer to what some authors call 'rule' commitment. The regulator cannot commit to the levels at which future standards will be set but he is (in the sense just outlined) able to commit himself to the <u>rule</u> he will follow in determining those levels[7].

In terms of the simple model we are working with 'rule' commitment implies that, in period 2, the owner of a plant licensed in period 1 will face regulatory standards determined by

$$(27)\quad \Omega.\beta(0) + [1-\Omega].\beta(r),$$

where $\beta(\bullet)$, recall, is the level of ALARA-compatible standards. The 'effective' level of regulations he faces is thus a linear combination of the 'raw' levels from the respective periods. The regulator is able to commit himself, by writing grandfather clauses into licensing agreements, to protecting existing licensees from some portion Ω of the regulatory impact of technological advances which occur during the lifetime of the total (i.e. the operator is required to meet just those regulations which were in force when his plant was licensed). When $\Omega=0$ there is no grandfathering and we revert to the case of ALARA

---

(7) This scenario is motivated by the position of the Nuclear Regulatory Commission (NRC) in the US, which can be seen as excercising discretion over the range of application and timing of execution of environmental standards which are legislatively mandated. In particular, as executive agency the NRC has discretion over 'compliance scheduling' - how rapidly plants already in operation should be obliged to comply with new regulations. Other environmental executive agencies (notably the US Environmental Protection Agency, (EPA)) find themselves in a similar position of 'constrained discretion'. Yaeger (1990) provides an insightful institutional examination of the scope of regulatory discretion excercised by the Environemntal Protection Agency in the US.

regulation where standards are driven by the costs and benefits of state of the art technology at all times[8].

The assumption that effective standards will be determined as a convex combination of raw standards is, of course, ad hoc. It can be defended, however, as a plausible interpretation of what legislators and others involved in policy generally mean by the term 'grandfathering' - that existing agents be protected from some portion of new requirements. It is also empirically justified as approximating the types of contractual clauses that are most often observed - there are likely to be legal and other institutional problems involved in drawing-up excessively obscure, non-linear regulatory contracts.

The Japanese nuclear industry has been granted licences containing various grandfather clauses over the years. Industry representatives in the US, lobby for the protection of existing plants against new regulations even more vociferously than for the regulatory relief of new plants. Raising finance for new projects is regarded as an almost impossible task given the

---

(8) A more thorough treatment or grandfathering would require a model having more than two periods that would more satifactorily capture the impact which grandfathering could be expected to have on the *intertemporal* choices of the firm. In such a model equilibrium would permit the possibility that neighbouring plants of different vintages could operate according to different standards. The implication would be that in the case of at least one of the plants the marginal cost of further increments of risk reduction would depart from the associated (myopic) marginal benefit. One would expect that for higher values of Ω the firm would wish to bring forward the construction of some its plants scheduled for later construction (to protect them from tightening of regulations in the meantime) and to delay some of its R&D expenditures until more of its plants were in place. Such considerations are outside the scope of our model.

aggressivity of the NRC's current back-fitting programme. Investors are unlikely to put up the billions of dollars of capital required for the construction of a new plant (which itself takes 8 to 10 years) without some commitment by the regulator regarding the regulatory standards which that plant will be obliged to satisfy 5, 10 or 30 years into its life. The NRC's aggressive stance on backfitting is a hot area of contention between the Commission and its subjects. At the time when the majority of todays plants were commissioned (the mid-1970's) those 'pioneer investors' did not forsee the spiralling increases in regulatory requirements, and consequently in operating costs, which were to follow. It is the 'experience' and innovative efforts of those investors which have generated the improvements in the safety of reactor designs to the point at which they are today.

The decision problem faced by the firm in the simple model we are using here is essentially static. The effect which we are particularly interested in is the role which grandfathering can play in encouraging more R&D effort by the industry. The industry is protected, to some extent, from rule changes induced as a result of the R&D efforts that they make. The regulator has discretion over the 'extent' of this protection because he chooses Ω. His rationale for choosing a Ω different from zero is that doing so may induce extra research effort from the industry.

Solving first the firm's problem over the choice of an R&D budget, given the grandfathering rule adopted by the regulator, the problem it faces and the associated first-order condition are

$$(28)\quad \text{Max}_r\{2b - [c(0,ß(0)) + c(r,\Omega.ß(0) + [1-\Omega].ß(r))] - r\}$$

$$(29)\quad -\partial P/\partial r = 1 + c_r(r^*, \Omega.ß(0) + [1-\Omega].ß(r^*))$$

$$+ c_S(r^*, \Omega.ß(0) + [1-\Omega].ß(r^*)).(1-\Omega).ß'(r^*)$$

$$= 1 + c_r + c_S.(1-\Omega).ß'(r^*)$$

$$= 0$$

In this case, then, the firm's first-order condition implicitly defines the choice of a level of safety related R&D expenditure as a function of Ω, the degree of grandfathering.

With Ω as his policy instrument, the regulator chooses the extent to which he 'grandfathers' regulations to maximise the expected value of the social welfare function (SWF), subject to the firm's incentive compatibility constraint. This, as noted earlier, can be seen as a stylisation of the problem faced by the Nuclear Regulatory Commission in the United States.

The planner's problem and associated FOC and SOC can be written;

$$(30)\quad \text{Maximise}_{\,0\le\Omega\le 1}$$

$$\{\mu.P(r^*) - d.[2 - ß(0) - \Omega.ß(0) - (1-\Omega).ß(r^*)]$$

$$(31)\quad \partial SWF/\partial\Omega = \mu.(\partial P(r^*)/\partial r).(\partial r^*/\partial\Omega) - \mu.c_s(r^*, \Omega.\text{ß}(0) + [1-\Omega].\text{ß}(r^*)).(\text{ß}(0) - \text{ß}(r^*))] - d.[\text{ß}(0) + \text{ß}(r^*)] + d.(1-\Omega).\text{ß}(r^*).(\partial r^*/\partial\Omega)] = 0$$

and

$$(32)\quad \partial^2 SWF(r^*)/\partial\Omega^2 < 0$$

These two equations, then, constitute the necessary conditions characterising the optimal choice of the degree to which regulatory standards should be retroactive. Solving them simultaneously will yield Ω^* - the optimal realisation of the policy instrument. That the expression is to be evaluated at r^* merely signals that adjustments which the firm may make to policy changes are taken into account. In more formal language the firm's reaction function, which constitutes the incentive compatibility constraint, has been built into the model. Therefore, $\partial P(r^*)/\partial r$ can be set equal to zero by application of the envelope theorem to the firm's problem. After some simplification the expression can be rearranged with the c_s partial on the left-hand side, making comparison to the pure ALARA case straight-forward;

$$(33)\quad c_s(r^*, \Omega.\text{ß}(0) + [1-\Omega].\text{ß}(r^*))$$

$$= d/\mu.\left\{1 - [(1-\Omega).ß(r^*)./[ß(r^*) - ß(0)]].\partial r^*/\partial\Omega\right\}$$

Recall that in the pure 2-period ALARA case, unadulterated by grandfather clauses, the regulator continually updates requirements to ensure that c_s is kept in equality with the constant (d/μ). Thus, maintaining the plausible assumption that c is increasing in s, optimal policy will be characterised by a value for Ω greater than zero when the ratio term in the curly brackets is negative (we will denote this optimal parameter value by Ω^*. It can be seen that the sign of this term will be the reverse of the sign of $(\partial r^*/\partial\Omega)$.

The result is analogous to one of those derived earlier. The planner wishes to encourage more research efforts viv-à-vis their ALARA-compatible levels. If $(\partial r^*/\partial\Omega)>0$ then grandfathering does encourage more R&D expenditure, the condition outlined in the last paragraph is satisfied and, therefore, $\Omega^*>0$. This case is illustrated in Figure 4.

The NRC abstains from discriminating between plants on the basis of vintage (ie sets $\Omega=0$). Industry representatives argue that there could be gains to everyone if the NRC were to use grandfather clauses more actively. It is not surprising that the industry would say this - expected producer surplus is everywhere increasing in Ω. In addition to the effects analysed here, it is also in the interest of existing producers to protect

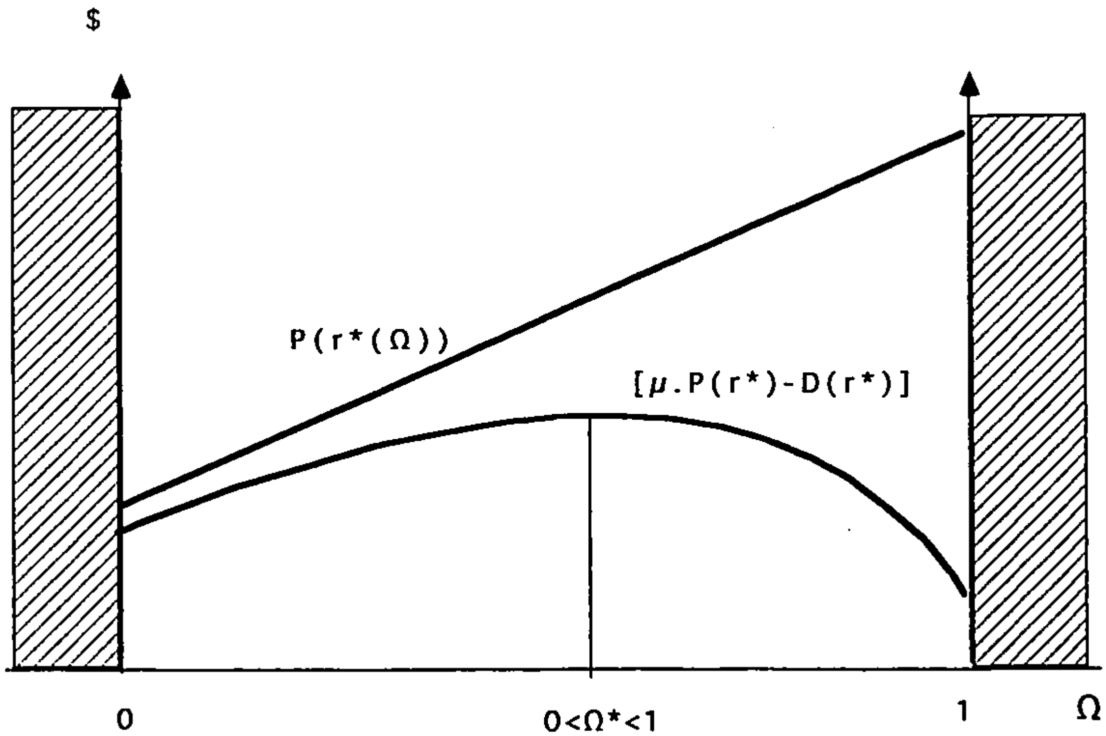

Figure 4 : The derived relationship between social welfare, industry surplus and the degree of grandfathering (Ω) under ALARA regulation, when condition given in text is satisfied..

the competetive position of old plants by insisting on tighter regulations for new plants. As can be inferred from the analysis in this section, however, when technological progress is endogenous optimal departures from ALARA can be expected to increase producer's surplus *and* decrease environmental risk contemporaneously.

The grandfather clause, $\Omega>0$, can be thought of as performing the function of a patent in other contexts. Once an invention has been developed the planner would prefer it to be exploited competetively, but is routinely willing to grant the inventor some monopoly power guaranteeing a patent for some fixed term in order to leave intact the dynamic incentives to invent. Grandfathering is analytically similar to issueing patents on safety technologies which the industry develops, where the choice of Ω is equivalent to the choice of the length of the patents' term. It is not surprising, then, that we have found that grandfather clauses may have agreeable normative properties (in comparison to the myopic application of ALARA rules) - the same agreeable properties commonly attributed to patents in other contexts.

(5) Dynamic Model of Safety with 'Learning-by-Doing'; On Learning to Operate a Nuclear Reactor Safely.

Not all innovation in an industry necessarily originates from the laboratories of the research and development division. As

Ashford and Heaton note, in the context of the chemical industry, "....innovation often occurs without R&D, prototypes or basic research, and is frequently the result of a trial-and-error process and experience", (Ashford and Heaton (1983: 111)). In this section we consider the possibility that operating a nuclear reactor or any other industrial installation in an environmentally friendly fashion may be something which is, loosely speaking, 'learned'. We thus employ a simple version of the type of learning curve analysed in a more general context by Spence (1980).

An extensive empirical literature points to the significance of learning by doing in the safe construction and maintenance of nuclear power plants (see, for instance, Zimmerman et al (1984) and the citations there-in). With such a complex and rapidly developing array of technology, and the potential for environmental damage they embody, it is hardly surprising that experience can vastly improve productivity both at the level of the individual employee, the operating unit and the power corporation (for examples of the types of mechanisms which may be at work in the context of this particular industry see Slater (1985)).

The analysis is similar to that used in the last sections, except that the cost function facing the industry will now have 'cumulative operating experience' as its first argument rather than 'cumulative R&D spending'. In this two period model the industry starts, by assumption, with zero experience. The cost function can be then be represented, for period 1 and 2 respectively; $c(0,s_1)$, $c(n,s_2)$, where n is the industry's (durable) output choice made at the start of the game. We will specify the cost function as generally as possible, surpressing other contributory factors, such that $c=c(n,s)$. It is plausible to suppose that $c_n<0$, $c_{nn}>0$, $c_s>0$, $c_{ss}>0$ and that $c_{ns}<0$. One can think of n as the number of plants, though we do not want to address the problems associated with the 'lumpiness' of investment opportunities in the electricity industry. Thus, we will assume formally that n measures the wattage of installed capacity and can vary smoothly between zero and infinity, thus maintaining the requisite differentiability of the industry's objective function. Diagrammatically we can think of the cost curve being convex in regulatory standards, with 'operating experience' generating downward movements in that curve in the usual space (Figure 5).

The movements illustrated imply, plausibly, that $c(n,s)$ approaches some finite lower bound as n increases without limit (for all $s<1$), and that $c_s(0,s) > c_s(n,s)$ for all s.

In the next sub-section we set about solving the industry's planning problem when it learns by doing and faces an ALARA motivated regulator.

(5:1) The Industry's Problem Under ALARA Regulation With Learning by Doing.

The profile of safety requirements set by the ALARA regulator

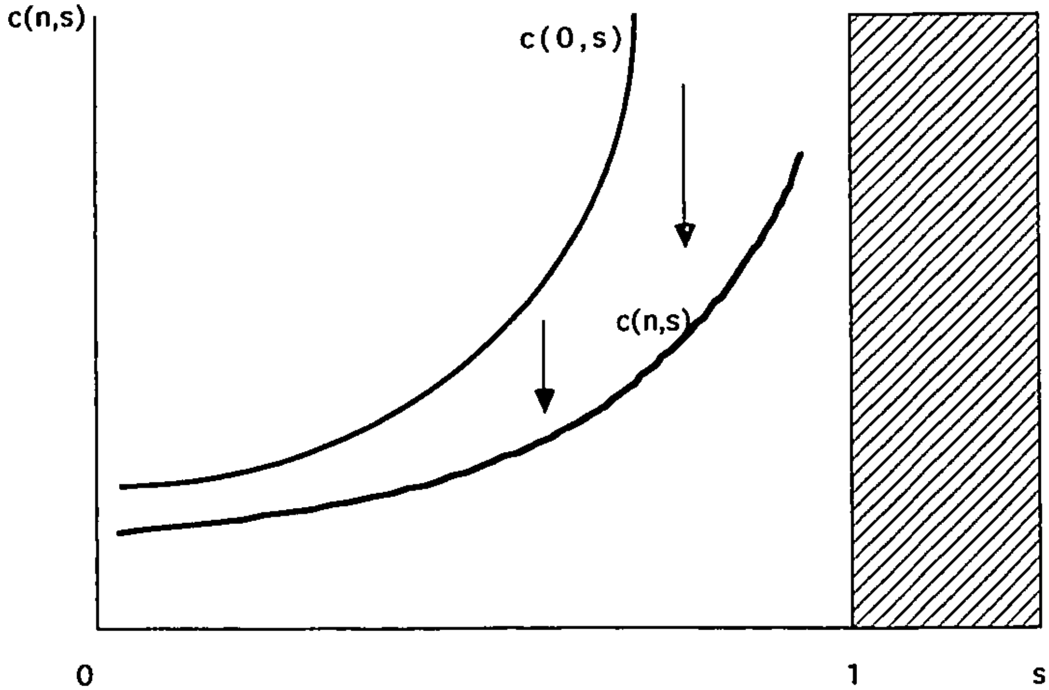

Figure 5: Impact of increase in first-period capacity on second-period cost of safety curve, due to learning by doing.

are easy to characterise in the 2-period case. 2nd period regulations are defined as an implicit function of the industry's first period choice of capacity. To maintain generality, we will suppose that standards at any time 't' are determined as a function of cumulative operating experience according to the function ¥. Application of the implicit function theorem reveals that $\partial ¥/\partial n=-\{c_{sn}/c_{ss}\}>0$. Thus, as operating experience increases the stringency of regulatory standards ratchets upwards correspondingly. ¥ in the learning-by-doing case corresponds to the ß function employed in earlier sections. Environmental standards dictated by the ALARA rule in the 2-period model which we have constructed will thus satisfy;

$$(34)\quad s_1=¥(0):\; c_s(0,s_1) = d/\mu$$

$$(35)\quad s_2=¥(n):\; c_s(n,s_2) = d/\mu$$

In the rest of this section we look at what the outcome looks like under ALARA. Further sections will address if and when the regulator could improve upon this outcome by systematically departing from the ALARA rule in setting required standards.

In the ALARA case the industry problem is to make the investment decision knowing that the standards will be updated, in the manner outlined above, at the start of the 2nd period. The industry problem and associated conditions are, therefore[9];

$$(36)\quad \text{Maximise}_n\ P=\left\{ 2.b.n\ -\ n.[c(0,s_1)\ +\ c(n,s_2)] \right\}$$

$$(37)\quad \partial P(n^*)/\partial n\ =\ 2.b\ -\ [c(0,s_1)\ +\ c(n^*,s_2)]$$
$$-\ n^*.[c_s(n^*,s_2).¥'(n^*)\ +\ c_n(n^*,s_2)]\ =\ 0$$

$$(38)\quad \partial^2P(n^*)/\partial n^2\ =\ -\ \{c_s.¥'(n^*)\ +\ 2.c_n$$
$$+\ n^*.¥'(n^*).[2.c_{sn}\ +\ c_{ss}.¥'(n^*)]\ +$$
$$n.c_{nn}.¥'(n^*)\ +\ n^*.c_s.¥''(n^*)\ \}$$
$$=J<0$$

The second-order condition for a maximum is assumed to hold (that is, $J<0$).

Once again, the asterisk denotes the solution to the industry problem. The FOC is composed of three composite terms. The first two represent the operating benefit from a marginal increase in capacity, net of any learning effects. The third term is the

---

(9) In setting-up the industry's problem in this case we abstract from the possibility (indeed probability) that the production of electricity will be characterised by other than constant returns to scale. Thus the benefit of operating an extra watt of generating capacity (net of environmental protection expenditures) is a constant. It would be more realistic to denote the net benefit of capacity by some non-linear function b(n), with $\partial(\partial b/\partial n)/\partial n>0(<0)$ constituting the case of increasing (decreasing) returns to scale in the production of electricity, but this would obscure our analysis whilst adding little or nothing of importance to our results.

interesting one from our perspective. It captures the impact which the extra period 1 learning associated with the marginal increase in capacity will have upon period 2 costs. This latter effect is itself comprised of two terms. The direct gains from having learned and thus become more efficient, and the impact which that learning will have upon costs through the increasedly regulations it induces. Each additional unit of capacity exerts, in the second period, an 'externality' on each existing plant equal to $[c_s(n^*,s_2).¥'(n^*)+c_n(n^*,s_2)]$ or, after some rearrangement,

$$(39)\ c_n.[1 + (ds/dn).c_s.c_n]$$

Notice that this externality may be positive or negative.

(5.2) Sub-Optimality of the ALARA Criterion in the Context of Learning-by-Doing.

To examine whether the industry can be expected to over- or under-invest vis-à-vis the socially preferred level of capacity, we evaluate the sign of $\partial SWF/\partial n$ at n^*. The planners objective is to maximise

$$(40)\ SWF = P - d.n.[2 - ¥(0) - ¥(n)]$$

Recognising that $(\partial P(n^*)/\partial n)=0$ by the envelope theorem;

$$(41)\quad \partial SWF(n^*)/\partial n = -d.\left\{[2 - ¥(0) - ¥(n^*)] - n^*.¥'(n^*)\right\}$$

$$= -d.\left\{2 - ¥(0) - ¥(n).[1+(n/¥).d¥/dn]\right\}$$

The final term in the lower expression is the <u>elasticity</u> of standards with respect to changes experience (ie accumulated 'doing'), which is positive. The sign of the whole expression is ambiguous. This implies that if there is learning-by-doing, employing the ALARA-algorithm to update regulatory standards may induce either excessive or inadequate investment in capacity during the 'learning phase'. Only in the rather improbable case in which $\partial SWF(n^*)/\partial n=0$ will use of the ALARA-rule constitute efficient policy.

If the elasticity term identified, $(n/¥).d¥/dn$, is sufficiently large (such that regulatory standards are comparatively sensitive to industry experience) then $\partial SWF(n^*)/\partial n>0$. In this case not enough capacity will be installed under ALARA regulation. The regulators rule obliges the industry to expend rather too much of their extra 'learning' on satisfying ever more challenging environmental requirements rather than on generating extra profits by meeting un- or little-changed standards at lower cost, and the industry is unwilling to make the requisite investment.

Conversely, if the elasticity term is comparatively small then ALARA regulation will encourage excessive levels of capacity vis-à-vis the social optimum. In this case too much of the social cost associated with the acquisition of experience is being borne by the public, in the form of too many dangerous period 1 plants.

Whichever of these two cases the parameters dictate should be of primary interest to the policy-maker in a particular context, we have demonstrated an *a priori* rationale for policy to depart from that dictated by the ALARA methodology. It is the form that this departure should take which we attempt to characterise during the rest of this section.

(5.3) The Paradoxical Possibility of Detrimental Learning Under ALARA Regulation.

In Appendix 1 we demonstrate that, when the regulator uses an ALARA rule, an exogeneous increase in the learning parameter (the steepness of the learning curve) may reduce social welfare. This is an unintuitive result and has important implications for policy.

In a recent paper, Lewis and Sappington (1991) develop a model in which exogeneous technological progress may be detrimental to social welfare. Whilst technological advances unambiguously serve to reduce innate production costs (an 'efficiency effect'), they also have a 'control effect'. That is to say they have an additional impact upon the nature of the solution to the principal-agency game played between a firm and it's sub-contractors, which will generally be welfare reducing. The overall social welfare implications of an increase in the rate of exogeneous technological advance is the sum of the 'efficiency' and 'control' effects, and is shown to be ambiguous. Technological advance may reduce social welfare - even when it arrives without (direct) cost as 'manna from heaven'.

In this contribution we demonstrate an analogous possibility (though with a quite different mechanism) in the context of an ALARA regulated industry which faces a 'learning curve' (see Spence (1981)) in the practice of externality control. With 'g' as the parameter capturing the 'rate of learning' (analogous to c_n, the reduction in unit cost resulting from an incremental unit of experience) we show that the relationship between g and social welfare is qualitatively ambiguous. The Appendix presents results for the general functional form. In addition, for concreteness we posit a simple logarithmic cost function and demonstrate that for plausible parameter values $\partial SWF/\partial g<0$ - an increased potential for learning-by-doing in the field of pollution or risk control reduces social welfare. Ceteris paribus, a shallow learning curve can be socially preferred to a steeper one.

As in Lewis and Sappington (1991), an increase in the parameter of interest (in our case g) has two counter-posed effects. Ceteris paribus, it reduces the period 2 costs of the firm of complying with a particular environmnetal standard (alternatively, it permits a cleaner period 2 environment without a reduction in producer's surplus). This 'direct effect' is unambiguously beneficial. The parameter change also, however, impacts upon the nature of the equilibrium of the standard-setting game played between the agency and the industry (in the ALARA case, of course, the agency's equilibrium strategy is to mechanistically update current period standarsd according to the ALARA-algorithm). When the learning curve is steep a small amount of learning in period 1 induces a relatively large increase in the stringency of requirements in the second period. The firm anticipates this and is thus less favourably disposed towards first period learning - a classic example of what some authors have termed (in other contexts) 'regulatory chill'. The welfare impact of this 'regulatory effect' is ambiguous, and may partially or fully offset the positive direct effect.

Another way of saying the same thing is that social benefit would, under ALARA regulation, be enhanced by permitting the industry to costlessly 'throw away' (forget), after period 1 is complete but before the regulations for period 2 are set, some part of what it had learned. Equivalently, if some mechanism existed whereby the industry had the option of not learning ('putting it's hands over it's eyes' in some sense) then it would often choose to exploit that option, and the resulting equilibrium would often socially dominate the pure-ALARA equilibrium.

(5.4) Optimal Standard-Setting by Stackelberg Leading Regulator With Commitment.

First we consider, in the new context of learning-by-doing, the case in which the regulator is able to commit credibly to a 2-period sequence of regulations. If the regulator chooses to commit himself in this way then the open- and closed-loop solutions to the game will coincide.

Facing a regulator who has committed himself to some profile of regulatory requirements, the industry problem and associated FOC become;

(42) $\text{Maximise}_n\ P = \left\{ 2.b.n - n.[c(0,s_1) + c(n,s_2)] \right\}$

(43) $\partial P(n^*)/\partial n =$

$$2.b - [c(0,s_1) + c(n^*,s_2)] - n^*.c_n(n^*,s_2) = 0$$

Equation 43 defines n implicitly as a function of s_1 and s_2:

(44) $n^* = f(s_1,s_2)$

(45) $\partial n^*/\partial s_1 = f_1(s_1,s_2) = c_s(0,s_1)/J < 0$

$$(46) \quad \partial n^*/\partial s_2 = f_2(s_1,s_2) = -\,[c_s(n^*,s_2) + n^*.c_{ns}(n^*,s_2)]\;/J$$

$$= -\,c_s.\left\{1 + (n/c_s).(\partial c_s/\partial n)\right\}/\,J$$

Thus the firm's choice of capacity is decreasing in the harshness of period 1 standards. It will be decreasing (increasing) in period 2 standards if the term in braces, in the final expression, is negative (positive). The algebraic term is the elasticity of marginal compliance costs with respect to changes in the number of units of capacity and is negative. A necessary and sufficient condition for $\partial n^*/\partial s_2 < 0$ is that the absolute value of this elasticity be greater than unity.

The planners problem is to choose a sequence of regulatory safety standards to maximise the SWF, subject (as usual) to the incentive constraint of the industry. This problem is, then;

$$(47) \quad \underset{\{s_1,s_2\}}{\text{Maximise}} \; \text{SWF} = \mu.P(n^*) - D$$

or, more explicitly,

$$(48) \quad \text{Max} \left\{\mu.f(s_1,s_2).[2.b - [c(0,s_1) + c(f(s_1,s_2),s_2)] - d.n.[2 - s_1 - s_2]\right\}$$

The pair of first-order conditions associated with an interior solution

$$(49)\quad \partial SWF/\partial s_1 = 0 = \Big\{ \mu.f_1.[2b - c(0,s_1) - c(n^*,s_2)] - \mu.f.[c_s(0,s_1) + c_n(n^*,s_2).f_1] - d.f_1.[2 - s_1 - s_2] + d.f \Big\}$$

$$(50)\quad \partial SWF/\partial s_2 = 0 = \Big\{ \mu.f_2.[2b - c(0,s_1) - c(n^*,s_2)] - \mu.f.[c_n(0,s_1).f_2 + c_s(n^*,s_2)] - d.f_2.[2 - s_1 - s_2] + d.f \Big\}$$

Some manipulation allows c_s to be isolated in each case; this makes the solution easy to compare with the ALARA case. Recall that following the ALARA rule implies that standards are continually updated to ensure that $c_s = d/\mu$. The necessary conditions characterising optimal policy in the presence of learning-by-doing reduce to;

$$(51)\quad s_1^*:\ c_s(0,s_1) = (d/\mu)\Big\{ 1 - f_1.[2 - s_1 - s_2]/f \Big\}$$

$$(52)\ s_2^*:\ c_s(n^*,s_2) = (d/\mu)\left\{1 - f_2.[2 - s_1 - s_2]/f\right\}$$

Recall that $f_1<0$ whilst the sign of f_2 was found to be ambiguous (depending upon whether the elasticity of marginal compliance costs with respect to changes in capacity, $(n/c_s).(\partial c_s/\partial n)$, was greater than or less than -1). Thus, if there is learning-by-doing, optimal regulation (if commitment is feasible) entails standards <u>tighter</u> than ALARA during the learning phase, whilst the standards in the mature phase (period 2) may be more or less stringent than those dictated by the ALARA formula. More concretely, period 2 standards will be more stringent iff the absolute value of the elasticity of marginal compliance costs with respect to changes in the number of units of capacity exceeds unity. Otherwise they will be set less stringently.

(5.5) Optimal Standard-Setting by Stackelberg Leading Regulator Without Commitment.

When the regulator is unable to commit credibly to the environmental standards that he will require of licensees in the future, sub-game perfectness of the closed-loop solution which we characterise requires that period 2 regulations conform to the ALARA criterion, that is;

$$(53)\ s_2=¥(n):\ c_s(n,s_2) = d/\mu$$

The firm's problem and associated FOC become, then;

$$(54)\quad \text{Maximise}_n\, P = \Big\{ 2.b.n \ - \ n.[c(0,s_1) \ + \ c(n,¥(n))] \Big\}$$

$$(55)\quad \partial P(n^*)/\partial n \ = \ 2.b \ - \ [c(0,s_1) + c(n^*,¥(n))]$$
$$- \ n^*.[c_s(n^*,s_2).¥'(n^*) \ + \ c_n(n^*,s_2)] \ = \ 0$$

The FOC implicitly defines the industry's choice of capacity, in this case, as a function of the standards set for period 1, where

$$(56)\quad \partial n^*/\partial s_1 \ = \ \{c_s(0,s_1)/J\} \ < \ 0.$$

Thus, when the planner is unable to commit to future standards, an increase in the stringency of period 1 environmental standards will induce a decrease in capacity chosen, and thus in the amount of learning that that the industry has achieved at the start of period 2. This in turn implies that period 2 standards will be laxer than they would otherwise be, since $¥'(n)>0$ for all n.

The requirement that the solution must be subgame perfect thus takes away s_2 as a (credible) policy instrument. Any sequence of regulatory standards $\{s_1,s_2\}$ in which s_2 is different from that dictated by the ALARA algorithm does not satisfy this refinement and will not satisfy our definition of an equilibrium[10].

The planner sets period 1 standards (the only policy variable over which he wields credible control) to maximise expected social welfare subject to the incentive constraint of the firm. This problem and the single associated first-order condition become, therefore;

$$(57)\quad \text{Maximise}_{s1} \quad n^*.\Big\{\mu.[2b - c(0,s_1) - c(n^*,¥(n^*))] - d.[2 - s_1 - ¥(n^*)]\Big\}$$

and

$$(58)\quad \partial SWF/\partial s_1 = \mu.(\partial n^*/\partial s_1).[2b - c(0,s_1) - c(n^*,¥(n^*)) - n^*.c_n(n^*,¥(n^*)) - n^*.c_s(n^*,¥(n^*)).¥'(n^*)] - \mu.n^*.c_s(0,s_1) + d.n^* - d.(\partial n^*/\partial s_1).[2 - s_1 - ¥(n^*) - n^*.¥'(n^*)] = 0$$

This expression can be rearranged with $c_s(0,s_1)$ isolated so that, again, we can compare it to the ALARA outcome. The 'optimal' sequence of regulatory standards in the no commitment case is

---

(10) Such an equilibrium would be 'dynamically inconsistent' in the language of Kydland and Prescott (1977), the classic paper on the subject.

$$(59)\quad c_s(0,s_1) =$$

$$(d/\mu).\left\{1 + (\partial n^*/\partial s_1).[¥'(n^*) - (1/n^*).(2 - s_1 - ¥(n^*))]\right\}$$

$$(60)\quad c_s(n,s_2) = (d/\mu)$$

During the learning phase (period 1) standards are set <u>more</u> stringently than they would be under ALARA regulation if and only if the expression in braces in the upper equality is greater than 1. A necessary and sufficient condition for this is that $\partial s_2/\partial n(n^*)$ be sufficiently small. Having arrived at the 'mature' phase (period 2), the regulator can do no better than revert to the myopic ALARA rule.

(5.6) Optimal Grandfathering of Standards in the Context of Learning-by-Doing.

In this section we return to the case in which the regulator is able to commit 'partially' in the sense outlined earlier. Thus period 1 standards are ALARA-compatible and effective period 2 standards may be grandfathered. Ω will again be used to represent the 'degree of grandfathering', such that effective period 2 standards are determined according to the linear combination of period 1 and period 2 'raw' standards, where Ω parameterises the respective weightings;

$$(61)\quad s_2 = \Omega.¥(0) + [1-\Omega].¥(n^*)$$

The industry problem (with the regulatory response built-in) and associated first-order condition are

$$(62)\quad \mathrm{Max}_n$$

$$n.\left\{2.b - [c(0,¥(0)) + c(n,[\Omega.¥(0) + (1-\Omega).¥(n^*)])]\right\}$$

$$(63)\quad \partial P(n^*)/\partial n = 2b - c(0,¥(0)) + c(n,s_2)$$
$$- n^*.[c_n(n,s_2) + c_s(n,s_2).[1 - \Omega].¥'(n^*) = 0$$

Equation 63 defines implicitly the industry's preferred capacity choice as a function of the degree of grandfathering, Ω, in the case in which the regulator is restricted to choose policy sequences which are temporally consistent. Application of the implicit function theorem to the firm's problem demonstrates that, ceteris paribus, $(\partial n^*/\partial\Omega)>0$ - an increase in grandfathering of regulations will increase the firm's period 1 capacity choice. When Ω is large the firm is less heavily 'punished', in terms of increased stringency of environmental rules, in period 2 for the learning it has invested in doing in period 1.

The planner's problem is to choose the optimal 'protection' to offer to the industry against escalations in the regulatory requirements which occur during the lifetime of the plant (ie to choose Ω^*). The policy problem is then;

$$(64) \quad \underset{0 \leq \Omega \leq 1}{\text{Maximise}} \left\{ \mu.P(n^*) - D \right\}$$

The associated FOC, $(\partial SWF/\partial\Omega)=0$ can be rearranged to yield the following summary expressions characterising the optimal regulatory sequence. Again c_s has been isolated to facilitate comparison with those standards dictated by the ALARA algorithm.

$$(65) \quad c_s(0,¥(0)) = (d/\mu)$$

$$(66) \quad c_s(n^*,\Omega.¥(0) + (1-\Omega).¥(n^*)) = c_s(n^*,s_2)$$

$$= (d/\mu).\left\{ 1 + Z \right\},$$

where

$$(67) \quad z = (\partial n^*/\partial\Omega).\Big[[1-\Omega].¥' - (1/n^*)[2-(1+\Omega).¥(0) - (1-\Omega).¥(n^*)] \Big] \Big/ \Big[¥(0) - ¥(n^*) \Big]$$

The sign of Z is ambiguous. Recall that $(\partial n^*/\partial\Omega)>0$, and that

¥(n*)>¥(0) for any positive value of n*. Thus the sign of Z is the opposite to that of the term in square brackets in the numerator. Letting Δ¥ denote [¥(n)-¥(0)] and Δn denote (n-0), such that [d¥/dn].Δn ≈ Δ¥, that term can be reexpressed as

$$(68)\quad (2/n).\{(1-\Omega).\Delta¥ - (1-¥(0))\}$$

This, evaluated around Ω=0, can be seen to be negative, implying that Z>0 and that the planner will prefer to stick with the corner solution - Ω*=0. That is, in the case of learning-by-doing a policy of unadulterated ALARA in both periods will socially dominate any policy involving grandfather clauses. One implication of this is that the ability to commit partially is of no value to the regulator in this case.

(6) A Model With R&D and Learning: The Case of 'Experience-Directed Research and Development' (EDR&D).

In this section we analyse policy-making when improvements in the technology used to control the externality come about through the combined forces of R&D investment and through learning-by-doing. In particular we consider the possibility of 'Experience-Directed Research and Development' (EDR&D), to be specified later, which we believe to be a good stylisation of the process which generates technical improvement in risk-reduction technologies in the civil nuclear power industry.

In this section we consider the case in which the firm chooses both a (durable) level of capacity and an R&D budget, taking into account the possible synergies which may exist between the two in a dynamic context.

When there are R&D and learning by doing (LBD) the operating costs of the firm at any given moment depend upon accumulated R&D, accumulated learning (for which a proxy is accumulated experience) and the safety-level to be achieved. The single-period costs of the firm will now be determined by the function c(r,n,s). In the context of the two-period model we are employing in this paper the sequence of costs faced by the firm in the two periods are, respectively, $c(0,0,s_1)$ and $c(r,n,s_2)$. As usual it is assumed that costs are convex in r, n and s; $c_r<0$, $c_{rr}>0$, $c_n<0$, $c_{nn}>0$, $c_s>0$, $c_{ss}>0$. We do not place restrictions on the signs of the cross-derivatives at this stage.

The most important role of operating experience in the development of new safety technologies in the nuclear industry is, however, that it tells the corporation how R&D effort should be directed. It is error, in the sense of plant behaving in unexpected ways, which helps the R&D experts to identify weak links in the technology employed and thus to better direct subsequent research effort. To capture this effect we assume that $c_{rn}<0$. The marginal productivity of research, $-(\partial c/\partial r)$, is increasing in operating experience. It is the 'amount' of experience which will determine how well targetted the research effort is.

(6.1) EDR&D and Nelson's View of the Role of Knowledge in Research Efficiency.

Various aut ͻrs have discussed the linkages between R&D effort, learning and production. In addition Nelson (1982) introduces a concept of 'knowledge' into a traditional sequential sampling type model of R&D. It may be possible to regard EDR&D as an extension of Nelsons contribution - providing a mechanism by which the stock of knowledge can be endogenised.

Nelson acknowledges that the idea of some concept of 'knowledge' often lurks behind discussion in this field - though never being adequately defined. "In many verbal discussions, R&D capabilities are proposed to be related to the strength of knowledge. While this proposition can be interpreted in a variety of ways, a promising one is lent by models that treat R&D as a search that can be . . . pointed in different directions" (Nelson (1982: 454)). While we do not employ a stochastic framework (and, indeed, regard the 'lumpy' view of technology implicit in search models as inappropriate in the context of nuclear safety research), our model does include some notion of the 'direction' of R&D.

Whilst not formalising where 'strength of knowledge' originates, Nelson discusses a number of interesting possibilities. He posits (tautologically) that ". . . strong knowledge is something which reduces the expected cost of any R&D achievement" (Nelson 1988:455). The concept could, for instance, be employed to permit meaningful catogorisation of scientific research into 'basic research' (investigation of underlying scientific principles) and 'applied research' (application of those principles) in terms of the extent to which a given project has cost externalities on downstream (read 'more applied') projects.

The interpretation which is closest to our's is Nelson's discussion of knowledge as something that serves to focus search. In the simplest case 'knowledge' tells the research programmer whether it is 'better' (in the sense of stochastic dominance) to sample from the set of 'blue projects' rather than 'yellow projects' in order to achieve a specified result. The hypothesis of 'Experience-Directed Research and Development' (EDR&D) is that it is operational experience which tells the research programmer which projects constitute the 'best bet' in reaching a given end.

The mechanism of EDR&D is only likely to be applicable in selected contexts. We believe that it provides an excellent stylisation of the process of technical advance in the field of risk-reduction in the nuclear power industry. The distinguishing characteristic of the industry is that scientists do not know the risks of malfunction associated with different parts of the hardware and therefore have only very hazy estimates of the marginal risk reduction contribution of improvements in various components. U.S.N.R.C. (1987) is the most comprehensive risk assessment of a national nuclear programme ever undertaken. It employs the event-tree methods of 'Probabilistic Risk Analysis' (an umbrella term for a set of risk assessment techniques). Yet the degrees of uncertainty surrounding risk estimates are enormous - in some cases 5% upper and lower significance bounds differ by factors of several thousand. Use of best estimates in such cases is of limited value, yet they are often quoted to several decimal places by the industry since, in general, they suggest that risks are small. Such practice may even be socially damaging, lending a pseudo-scientific air of certainty to a situation where, in actual fact, very little is known.

The most important source of information which can narrow the confidence limits within which plant safety can be assessed is operating experience. The basic safety philosophy employed in Western-designed power plants is that of 'defence in depth'. The associated engineering methodology is that the nuclear plant should incorporate a sequence of functionally independent safety mechanisms or bulwarks, each individually capable of averting a plant accident, such that each must fail in sequence if an accident is to ensue. If enough such safety 'layers' are put in place, and their probabilities of malfunctioning at the critical time are genuinely independent, then the probability of an accident occurring can be made vanishingly small. The designer can always add another tier of defence (inappropriately termed 'redundant defences') - at a cost - to reduce further the compound probability associated with an accident.

The safety tiers in a typical U.S. light-water reactor (LWR) include various engineered mechanisms capable of 'scramming' (closing down) the reactor, cooling systems capable of slowing down the rate at which a malfunctioning reactor core heats up and a series of inert containtment structures. Operating experience yields statistical information to the industry regarding the reliability of the various systems. In the case of an inert system this will be in terms of failures per call, in the case of an engineered system in terms of percentage down-time. A manufacturer may have estimated, a priori, that a given system would be available 92% of the time but long-run statistical evidence from sufficient operating experience may cause this estimate to be revised to 85%. This may serve to take the plant below the threshold of acceptable risk and induce the regulator to order back-fitting of additional redundant systems. More important for the purposes of this paper, this type of information, collected in the operating logs for every component, will serve as a primary input for identifying weak links in the reactor design and thus to targetting subsequent research effort.

The basis of 'Experience-Directed Research and Development' (EDR&D) is that operating experience permits 'targetting' of R&D effort. Many other authors have discussed the interaction between R&D and learning.

Foray (1991) notes that the two aspects of technical advance, learning and R&D, are inextricably linked. He uses this to develop hypotheses about the non-equivalence of R&D conducted inside rather than outside the firm; "Il ne peut y avoir , contrairement à ce que soutenait implicitement Arrow, de substituabilité parfaite entre une ressource de R&D interne et une ressource de R&D externe. Il convient d'insister plutôt sur les aspects de complémentarité liés à l'irréductibilité de la fonction d'apprentissage, assumé par la R&D interne", (Foray (1991: 785)). In particular he identifies complementarities between learning by doing and formal research[11].

The recognition that the action of R&D may impact upon the learning capability of the firm is, as these authors have demonstrated, of considerable importance. It is also possible that the action of R&D may impact upon the R&D abilities of the firm (ie the activity of research may itself be subject to learning-by-doing) and/or the action of learning may impact upon the learning abilities of the firm (ie the learning curve may be convex in the usual space). Such cases have often been considered in both the theoretical and empirical literature on the innovation process. The novel contribution of this section is to have

---

(11) Cohen and Levinthal (1989) support the view presented by Foray that the act of performing research and development may in itself improve the firms ability to learn; ". . . we argue that while R&D obviously generates innovations, it also develops the firms 'learning' capacity", (Cohen and Levinthal (1989: 569)). They go on to note that the impact which R&D may have upon learning has been neglected by students in the field (Allen (1977) and Mowery (1983) are exceptions). In addition, Zhou and Rubenstein (1986) develop a concept of 'Imbedded Technology Capability' (ITC) in the context of science and technology policy in Chinese manufacturing industry, which has similar characteristics. They pay specific attention to the interface between R&D effort and production, noting that 'know how' accumulated by the firm as a by-product of R&D can be expected to increase the subsequent capability of the firm to learn.

identified a mechanism whereby operating experience impacts upon the R&D abilities of the firm.

(6.1) Equilibrium Under ALARA Regulation with EDR&D.

As in earlier sections, ALARA regulation implies that environmental standards are continually updated to ensure equality between the marginal costs and benefits of additional increments of risk reduction *given* the technology in existence. In the two-period case the sequence of regulatory standards that the industry will be obliged to satisfy, $\{s_1, s_2\}$, will satisfy the following pair of equalities;

$$(70)\quad c_s(0,0,s_1) = d/\mu = c_s(r,n,s_2)$$

where r and n are the realised levels of R&D expenditure and operating capacity. ALARA standards are, then, implicitly defined as a function of accumulated R&D expenditure *and* accumulated operating experience. Let Ø represent the function thus generated such that, in the two-period context s_1=Ø(0,0) and s_2=Ø(r,n).

The implicit function theorem yields the derivative properties of the function Ø, in particular $Ø_r(r,n)=-c_{sr}(r,n)/c_{ss}(r,n)$, which is positive and $Ø_n(r,n)=-c_{sn}(r,n)/c_{ss}(r,n)$, which is also positive.

It is straight-forward to solve the industry's problem when it is faced with an ALARA-motivated regulator. The firm moves to maximise the two-period sum of producer's surplus (as in earlier sections we abstract from discounting since we are not interested in effects generated by time-preference) by choosing a level of capacity and an R&D budget;

$$(71)\ \underset{\{r,n\}}{\text{Maximise}}\ P= \left\{2.b.n - n.[c(0,0,s_1) + c(r,n,s_2)] - r\right\}$$

subject to $s_1=Ø(0,0)$, $s_2=Ø(r,n)$

The pair of first-order conditions which characterise the solution to the firm's problem, $\{r^*,n^*\}$, are then;

$$(72)\ \partial P(r^*,n)/\partial n = 0$$

$$= \left\{2b - \left[c(0,0,Ø(0,0)) - c(r^*,n,Ø(r^*,0))\right]\right.$$

$$\left. - n.\left[c_n(r^*,n,Ø(r^*,n)) - c_s(r^*,n,Ø(r^*,n)).Ø_n(r^*,n)\right]\right\}$$

and

$$(73)\ \partial P(r,n^*)/\partial r = 0 =$$

$$- n^*.\left[c_r(0,0,Ø(0,0)) - c_s(r,n^*,Ø(r,n^*)).Ø_r(r,n^*)\right] - 1,$$

corresponding to the equations numbered (37) and (12) respectively. The various components in the FOC 's have familiar interpretations. In each case the firm increases the level of its choice variable up to the point of equating the incremental private cost and benefit. These include the dynamic benefits associated with technological improvements as well as the costs associated with the tightening of regulatory standards induced by those technological improvements.

(6.2) The Welfare Sub-Optimality of the ALARA Rule Under EDR&D.

It is of significant policy interest to demonstrate that, in the general case, the presence of EDR&D will serve to negate the welfare optimality credentials of ALARA regulation derived in section 3. Doing so provides a rationale for considering other policy rules.

In this sub-section we characterise the inefficiency of the ALARA rule as a basis for standard-setting. In doing so it will be possible to consider both the quantitative and *qualitative* characteristics of the research and development programme. The safety-R&D budget is of dollar value r. The level of capacity in the first period will dictate the amount of 'learning' (in the sense of 'learning how to target research money', as discussed earlier in the section) which occurs, and can therefore be seen as an index of how well the R&D funds are directed. Thus n, the capacity choice, can also be interpreted as a proxy for the 'quality' of the safety-R&D programme, in the sense that it is n which determines the marginal productivity of each dollar of the R&D budget.

In Section (6.3) we attempt to characterise what optimal policy *will* look like, and how it differs from the ALARA case. Unlike for R&D and LBD, in the context of EDR&D we only investigate optimal policy formulation in the case of a Stackelberg leading regulatory bureau with commitment.

Recall that the planner seeks to maximise a social welfare function (SWF) which is a weighted sum of producer's surplus and the expected value of environmental/public health risk:

$$(74)\quad SWF(r,n) = \mu.P - d.n.[2 - s_1 - s_2]$$

To demonstrate the sub-optimal characteristics of ALARA-motivated standard-setting in the presence of EDR&D we consider the two relevant partial derivatives in the vicinity of the equilibrium induced by ALARA rules, derived in Section (6.1) above.

$$(75)\quad \partial SWF(r^*,n^*)/\partial r = \mu.\partial P(r^*,n^*)/\partial r + d.n^*\O_r(r^*,n^*)$$

$$= d.n^*\O_r(r^*,n^*) > 0$$

$$(76)\quad \partial SWF(r^*,n^*)/\partial n = \mu.\partial P(r^*,n^*)/\partial n - d.[2 - \emptyset(0,0) - \emptyset(r^*,n^*)] + d.n^*\emptyset_n(r^*,n^*)$$

$$= d.n^*.\emptyset_n(r^*,n^*) - d.[2 - \emptyset(0,0) - \emptyset(r^*,n^*)]$$

In each case the envelope theorem has been applied to the firm's problem.

The determination that $\partial SWF(r^*,n^*)/\partial r$ is unambiguously positive (where, recall, the asterisks denote that the derivative has been signed in the neighbourhood of the ALARA solution) implies that ALARA will always induce the firm to spend too little on R&D aimed at improving environmental protection technology. The sign of $\partial SWF(r^*,n^*)/\partial n$ is, however, ambiguous. This implies that, under an ALARA régime, the safety R&D budget may be more or less accurately targetted than would be socially optimal. It will be under- (over-) targetted if $\emptyset_n(r^*,n^*)$ is sufficiently large (small). Since $\emptyset_n(r^*,n^*) = -c_{sn}(r^*,n^*)/c_{ss}(r^*,n^*)$, which is greater than zero, we can restate this condition as implying that the R&D effort will be under- (over-) targetted if is sufficiently large (small) in absolute value. That is, when an increase in experience causes a sufficiently significant reduction in the marginal cost, to the regulator, of additional increases in regulatory stringency.

This is an interesting result. If the regulatory bureau chooses to update environmental standards in accordance with the ALARA algorithm the industry's R&D efforts will be quantitatively inadequate, but the effort which is expended may be more or less well directed than is optimal depending upon the size of $(\partial^2 c/\partial s \partial n)$. The possibility that it be better targetted than socially optimal follows from the recognition that 'targetting' has a social cost (borne externally to the firm) in the form of relatively high-risk period 1 plants. This cost is associated with the square bracket term in the expression for $\partial SWF(r^*,n^*)/\partial n$ above. As this externality becomes larger it becomes more likely that the R&D will be over-targetted. The firms recognition that rules will be ratchetted upwards as R&D becomes better directed and, thus, its results greater, serves as the restraining influence on the firm. If this restraint is sufficiently great ($\emptyset_n$ sufficiently large) then R&D becomes under-targetted. Nothing can be said, a priori, about which of these cases is more likely. This will depend upon the functional form of c(r,n,s) adopted.

(6.3) Optimal Standard-Setting by a Stackelberg Leading Regulator Without Commitment In the Presence of EDR&B.

In this section we consider briefly what the optimal policy rule is in the context of experience-directed research and development. The treatment is less comprehensive than in the contexts of R&D and 'classical' learning by doing considered in

Sections (4) and (5) respectively. In particular, we only solve the closed-loop problem in which the regulatory agency moves first and is unable to commit credibly to second period requirements at the start of the two period game. This inability of the agency to 'tie it's own hands' means that we restrict our attention to that sub-set of candidate solutions which are sub-game perfect ('dynamically consistent').

The agency's problem is to choose a period one safety requirement, s_1. Dynamic consistency implies that the regulator will revert to the ALARA algorithm in period 2 (ie $s_2=\emptyset(r,n)$). The agency, by assumption, moves first. Dynamic programming dictates that we solve the follower's problem first.

The firms problem is to determine both a level of capacity and an R&D budget to maximise its lifetime (ie two-period) producer surplus. Again, we restrict our attention to those (plausible) cases in which producer surplus is maximised by an interior value for both of the choice variables. It is straight-forward to show that the two FOC's associated with the firm's problem reduce to;

$$(77)\quad \partial P/\partial r = -n.[c_r(r,n,s_2) + c_s(r,n,s_2).(\partial s_2/\partial r)] - 1 = 0$$

and

$$(78)\quad \partial P/\partial n = 2.b - [c(0,0,s_1) + c(r,n,s_2)]$$
$$- n.[c_n(r,n,s_2) + c_s(r,n,s_2).(\partial s_2/\partial n)] = 0$$

or, recognising that s_2 will be determined by the ALARA rule

$$(79)\quad \partial P/\partial r = -n.[c_r(r,n,\emptyset(r,n)) + c_s(r,n,\emptyset(r,n)).\emptyset_r] - 1 = 0$$

and

$$(80)\quad \partial P/\partial n = 2.b - [c(0,0,s_1) + c(r,n,\emptyset(r,n))] -$$
$$n.[c_n(r,n,\emptyset(r,n)) + c_s(r,n,\emptyset(r,n)).\emptyset_n] = 0$$

To reduce arithmetic we will assume that the third derivatives of the cost function are negligible (zero).

Of interest to us are the changes in industry behaviour induced by marginal changes in the parameter over which the policymaker has credible discretion, namely s_1. Cramers rule dictates that

$$(81)\quad dr^*/ds_1 = [- P_{nn}.P_{rs} + P_{ns}.P_{rn}]/|\Delta| = -c_s(0,0,s_1).P_{rn}/|\Delta|$$

$$(82)\quad dn^*/ds_1 = [- P_{ns}.P_{rr} + P_{rs}.P_{rn}]/|\Delta| = c_s(0,0,s_1).P_{rr}/|\Delta|,$$

where

$$|\Delta| = P_{nn}.P_{rr} - (P_{nr})^2$$

is positive by virtue of the sufficiency conditions associated with the firms problem. Thus, dn^*/ds_1 is negative, whilst the sign of dr^*/ds_1 is ambiguous and is the opposite of the sign of P_{rn}.

The planner chooses a period 1 standard to maximise lifetime social welfare, subject to the incentive constraints which characterise the industry's reaction to policy reform (as captured by the derivatives in Equations 81 and 82). His supposed inability to commit credibly to period 2 standards means that standards in that second period will inevitably be ALARA-determined. His problem, then, is to choose s_1 such as to

$$\text{(83)} \quad \text{Maximise} \left\{ \mu.P - d.n.(2 - s_1 - \emptyset(r,n)) \right\}$$

The planners single first-order condition, recognising that application of the envelope theorem to the firm's problem implies that $(\partial P/\partial n).(dn^*/ds_1)$ and $(\partial P/\partial r).(dr^*/ds_1)$ can both be set equal to zero, is

$$\text{(84)} \quad \partial SWF/\partial s_1 = \mu.\partial P/\partial s_1 + (dr^*/ds_1).d.n.\emptyset_r - (dn^*/ds_1).d.((2-s_1-\emptyset) + n.\emptyset_n) + d.n = 0$$

Which, after substituting for the three partial derivative terms becomes

$$(85)\ \partial SWF/\partial s_1 = 0 =$$
$$d.n - \mu.n.c_s(0,0,s_1) - (c_s(0,0,s_1).P_{rn}/|\Delta|).d.n.\text{Ø}_r$$
$$- (c_s(0,0,s_1).P_{rr}/|\Delta|).d.((2-s_1-\text{Ø}) + n.\text{Ø}_n)$$

As in earlier sections we re-express this condition (which characterises optimal policy in the context of EDR&D), isolating $c_s(0,0,s_1)$ in order to facilitate comparison with the pure ALARA case. We note that P_{rn}, P_{rr} and $|\Delta|$ are all insensitive to changes in $c_s(0,0,s_1)$, d and μ, such that they can legitimately be left unevaluated during this rearrangement. If the regulator used the ALARA rule this would imply that he would ensure that $c_s(0,0,s_1)$ equalled (d/μ). Rearrangement of Equation 85 implies that optimal policy will require, rather, that

$$(86)\ c_s(0,0,s_1) = (d/\mu).\left\{1/[n^2 + n.(d/\mu).M]\right\}$$

where

$$(87)\ M = (1/|\Delta|).\{P_{rr}.((2-s_1-\text{Ø}) + n.\text{Ø}_n) + P_{rn}.n.\text{Ø}_r\}$$

Consider the denominator of the term in braces in Equation

86:

(88) $[n^2 + n.(d/\mu).M]$

That the firms costs are convex in s implies that a necessary and sufficient condition for optimal regulatory standards to be more (less) stringent than ALARA is that this term be less than (greater than) unity. The relative size of tyhis expression cannot be determined a priori, and thence the policy implications of EDR&D in the regulated industry are qualitatively uncertain.

The ambiguity, whilst analytically disappointing, is not surprising. It reflects the complexity of the interactions between the three strategic variables in our analysis. When the technologies of environmental protection are subject to EDR&D, it was shown that the ALARA-regulated industry will not do enough R&D (ie r will be too small) whilst whether it will do too much or too little learning is uncertain (ie n may be too big or too small). In addition, it was shown that regulatory relief (ie committing to period one standards less stringent than their ALARA counterparts) would induce the regulated firm to reduce the amount of learning it does (ie $dn^*/ds_1<0$) whereas its impact upon the firms R&D budget was ambiguous (dr^*/ds_1 >/< 0: Equation 81). The likely normative impacts of regulatory relief (or, analogously, increased regulatory stringency) are, then, highly ambiguous and depend upon the various first- and second-derivative properties of the cost function c(r,n,s).

Simplistic rules which make generalised predictions as to how far an industry such as the nuclear sector should be offered regulatory relief whilst their safety technologies remain immature are likely to be just that - simplistic. Generalised answers to essentially empirical questions are unlikely to be of much use, we have shown that to draw even qualitative conclusions requires <u>empirical</u> estimation of c(r,n,s).

(7) Brief Conclusions - Reactive Regulation and the Sources of Technical Knowledge.

The aim of this paper was to dispel the notion, increasingly popular in the environmental regulation literature, that the ALARA criterion is an appropriate basis for standard-setting. We chose the nuclear industry for illustrational purposes, but the lessons learned are clearly generaliseable.

ALARA is essentially a static concept, atleast in the sense that its agreeable normative properties are demonstrable in a static setting. In a dynamic context it will in almost all cases be dominated by some other policy rule. When the flow of technical know-how regarding safety technologies is endogenous it is essential that the algorithm chosen to update regulatory standards through time be dynamically as well as statically optimal. In particular, it must leave enough of the surplus generated by technical advance as producer surplus, in order to sustain the industry's incentives to furnish an adequate, ongoing supply of new techniques. Only in extremely contrived circumstances will the ALARA algorithm be fortuitous enough to achieve this. More generally the policy-maker should depart from the ALARA rule.

The thrust of our results may appear to be destructive. A generalised algorithm for policy prescription, the ALARA rule, has been criticised and replaced with a 'general theory' which generates highly ambiguous, context-specific predictions. The context-specificity of the results reflects the context-specificty of optimal policy response and cannot be ignored by the policy analyst merely in order to simplify his conclusions. We hope to have demonstrated that the supremacy of the ALARA methodology is not robust theoretically. Furthermore, the sub-optimality of the practice is likely to be particularly acute in industries where the potential for technical advance in the field of environmental protection technology is great, such as the nuclear energy sector. This makes it of particular concern that ALARA is coming to be popularly exulted as the the 'way ahead' in the environmental regulation of that sector. The optimality credentials of the ALARA algorithm are generated in static settings, but are unlikely to be repeated in any but the most trivial dynamic setting.

Secondly, we have embarked upon the programme now necessary - that of characterising optimal departures from ALARA. We have, in particular, noted the qualitative distinctions between how regulation should be conducted in the context of an industry in which technology advances through research effort, versus one in which it advances primarily through learning by doing.

The key policy implications of our analysis are as follows. Standards compatible with the ALARA rule are, in general, not optimal in a dynamic context. In the R&D case the regulator can generate welfare improvment by adjusting period 2 standards if and only if he can commit to those adjustments at the start of the game. In the LBD case he can generate improvements by manipulating both period 1 and period 2 standards simultaneously, whether or not he can commit in this way. Requiring subgame perfection of the regulatory strategy (i.e. robbing the regulator of his ability to commit) thus precludes the possibility of welfare improving departures from ALARA in the former case (R&D) but not the latter (LBD).

Appendix.

ALARA Regulations and the Possibility that Learning Will Be Detrimental to Social Welfare.

In a recent paper Lewis and Sappington (1991) develop a model in which technological progress may be detrimental to social welfare. Whilst technological advances unambiguously serve to reduce innate production costs (an **'efficiency effect'**), they also have a **'control effect'**. That is to say they also have an impact upon the welfare characteristics of the outcome of the principal-agency game played between a firm and its sub-contractors. The overall social welfare implications of an exogenous technological advance is the sum of the 'production' and 'control' effects, and is found to be ambiguous. Technological advance may reduce social welfare.

In this note we wish to demonstrate a similar possibility in the context of an ALARA regulated industry which faces a 'learning curve' (see Spence (1981)) in the practice of externality-control. With 'g' as the parameter capturing the 'rate of learning' (assumed to be a constant and corresponding to the derivative c_n in the main body of the paper) we derive optimal firm behaviour under an ALARA régime. Further, it is shown that the relationship between 'g' and social welfare is qualitatively ambiguous. For concreteness we posit a simple logarithmic cost function and demonstrate that for plausible parameter values

$\partial SWF/\partial g<0$ - an increased potential for learning-by-doing in the field of pollution or risk control reduces social welfare. Ceteris paribus, a shallow learning curve could be socially preferred to a steeper one.

As in Lewis and Sappingtons' paper, an increase in the relevant parameter (in our case g) has two counter-posed effects. Ceteris paribus, it reduces the period 2 costs of the firm of complying with a particular environmental standard (alternatively, it permits a cleaner period 2 environment without a reduction in producer surplus). This 'direct effect' is unambiguously beneficial. The parameter change also, however, impacts upon the equilibrium of the standard-setting game played between the agency and the industry (in the ALARA case, of course, the agency's equilibrium strategy is to mechanistically update current period standards according to the ALARA-algorithm). The welfare impact of this 'regulatory effect' is, in general, ambiguous and may partially or fully offset the positive direct effect.

The Steepness of the Learning Curve and Social Welfare: Some Ambiguous Results.

[Case 1]: Generalised Functional Form.

Maintaining the notation used in the main text, let g be some parameter capturing the 'rate-of-learning' (the steepness of the learning curve). Increases in the value of g represent increases in the rate of learning-by-doing (LBD). $g \geq 0$ (we do not consider the possibility of 'negative' learning), where $g=0$ implies that the learning curve is horizontal (that is, there is no LBD).

Fully differentiating the additive SWF used in the paper (assuming, for brevity of notation, zero discounting), and substituting from the results in the text, it is straightforward to calculate that

$$(A1)\ dSWF/dg =$$
$$\mu.[2b - c_n.n - ¥'(n).c_s.n - c(0,¥(0)) - c(n,¥(n))].(\partial n/\partial g)$$
$$- \mu.(\partial c/\partial g) - (\partial n/\partial g).d.[2 - ¥(0) - ¥(n) - ¥'(n).d.n].$$

Evaluating this in the neighbourhood of equilibrium, applying the envelope theorem to the firms capacity-choice problem and simple rearrangement yields

$$(A2)\ dSWF(n^*)/dg$$
$$= - \mu.(\partial c(n^*,¥(n^*))/\partial g)$$
$$- (\partial n^*/\partial g).d.[(2 - ¥(0) - ¥(n^*)) - ¥'(n^*).d.n^*].$$

Thus, in the general case the sign describing the marginal general equilibrium impact of an increase in the steepness of the learning curve upon social welfare is ambiguous. The first composite term is the the 'direct effect' the second is the 'regulatory effect'. If the direct effect of an increase in g upon industry costs is, in the vicinity of equilibrium, sufficiently large (that is $\partial c(n^*,¥(n^*))/\partial g$ is a large enough negative number) and/or that cost reduction is weighed sufficiently heavily in the evaluaion of social welfare (μ is large) then the socially beneficial direct effect will outweigh the regulatory effect.

[Case 2]: Logarithmic Learning.

In this section we posit a simple functional form for the cost-function and demonstrate the possibility of 'perverse' results of the type discussed earlier in this Appendix.

Suppose that the cost function could be represented

$$(A3)\quad c(n,s) = e^{-gn+s}$$

where g characterises the rate at which costs decrease with experience and is scaled such as to lie on the unit interval. Notice that $c_n<0$, $c_{nn}>0$, $c_s>0$, $c_{ss}>0$ (as assumed in the text). The important point is that $c_{ng}<0$ - an increase in the parameter g increases the rate at which accumulated experience generates cost decreases, ceteris paribus.

If the regulatory régime is ALARA-motivated then the sequence of regulatory standards can readily be shown to be;

$$(A4)\quad s_1 = \log(d/\mu)$$

$$(A5)\quad s_2 = \log(d/\mu) + g.n$$

where all logarithms are to the the natural base. Notice that the stringency of period 2 environmental standards is linearly increasing in first period capacity (more concretely, $(\partial s_2/\partial n)=g$) - this linkage is the source of all of the incentive problems which this paper investigates. Application of the envelope theorem to the industry problem and a series of straightforward substitutions and rearrangements yields the following expression for the full-derivative of interest, evaluated in the vicinity of equilibrium;

(A6) $dSWF(n^*)/dg =$

$$\mu.n^{*2}.e^{-gn^*+s} - (2n^*d/g)(1 - \log(d/\mu)) > \text{ or } < 0$$

In equilibrium the qualitative impact of a steepening of the learning curve upon social welfare is ambiguous. The 'perverse' case will hold (ie $dSWF(n^*)/dg$ will be negative) when $(1 - \log(d/\mu))$ is a sufficiently large positive number. This is most likely to occur when d is small and/or μ is large.

Conclusions to Appendix.

In this note we have demonstrated the possibility that, when ALARA regulation is used (as, for instance, it is in environmental regulation in the United Kingdom and Italy, amongst other places) a steep learning curve in the domain of externality-control may be detrimental to social welfare. This was done both in the general case and in the case where costs are assumed to decline logarithmically with accumulated experience.

This is an apparently perverse result which follows from taking explicit account of the impact which a perturbation of the learning parameter may have upon the characteristics of the solution of the standard-setting game played between the regulatory agency and the regulated industry. The firm dislikes the large upward revisions in regulatory stringency induced by learning. When the learning returns to capacity are unusually large then the firm will reduce learning by reducing it's 'doing'. It is the social costs of this reduction that may induce welfare to be lower when learning is rapid. The result is reminiscent of a recent result of Lewis and Sappington, as outlined.

Market Structure, Regulatory Externalities and the Overloading of Trucks: An Industry Perspective.

(1) Introduction to the Problem.

The overweighting of trucks can do considerable damage to the man-made environment. By causing degradation to the pavement surface, overladen trucks impose costs on other road-users in terms of increased travel times and decreased comfort and road-safety. By increasing the amount of money which must be spent on road repairs the behaviour also imposes a burden on the public purse and hence upon society in general.

The economic incentive for an operator to overload his truck is clear - his revenue will increase approximately linearly in the weight of produce which he hauls, whilst his operating costs increase very little. Much of the extra cost of excessive loading is, then, external to the operator, taking the form of increased infrastructural damage. Because of the externality imposed by trucks which are overloaded, weight limits are set. The enforcement component of regulation is of unusual importance, in this context, because of the potential for 'outliers' to do disproportionate amounts of damage. A 'rule of thumb' commonly used by civil engineers is that the severity of the pavement damage caused by the transit of a truck increases with the fourth power of gross axle weight (Stiglitz and Arnott (1988:33)); one severely overweighted truck can do the damage of hundreds of compliant ones.

One study estimates that overloaded trucks in Saskatchewan (which constitute 13% of loaded trucks and, therefore, significantly less than 13% of total truck movements in that province) are responsible for 98% of all damage to highways (Wyatt and Hassan (1984: 70)). Thus, according to Saskatchewan Highways and Transportation, the practice of overweighting accounted for some Can$ 2 million of pavement damage in that province in 1984. Barron (1985) calculates that pavement damage attributable to overweight trucks in New Jersey could exceed US$ 20 million per annum. Similar conclusions regarding the disproportionate impacts of violators of weight regulations have been reached for other jurisdictions. It should be noted that these estimates often understate the true damage done to infrastructure because studies (including that by Wyatt and Hussan (1984)) ignore damage done to bridges and other roadside structures. More general estimates of the potential effects of heavy trucks have been provided by the U.S. General Accounting Office (1979) and the Organisation for Economic Cooperation and Development, Environment and Transport Divisions (OECD, 1983). What is clear is that the damage associated with the antisocial activity of overloading lorries is considerable.

In this paper we assess the regulatory problem surrounding the enforcement of weight limits. In particular, the scope for industry self-regulation is considered.

Regulation consists of a weight limit and a level of enforcement effort which is assumed to depend upon the compliance performance of the industry. Weight regulations vary between provinces in Canada, and between States in the US. Maximum loads are defined for each axle and for gross combination weight (GCW) by class of truck. Setting standards is not enough, in itself, however, to protect the road system from excessively heavy lorries. It cannot be assumed that operators will be uniformly scrupulous in complying with the regulatory requirements unless an adequate enforcement régime is installed. In this paper we concentrate on the problem of enforcing existing regulations, rather than that of at what stringency regulations should be set in the first place. In the context of a 'responsive' regulatory agency, we demonstrate that there are potential gains to be reaped by the industry from establishing a <u>self</u>-regulatory system to ensure a higher rate of compliance than would be the case if each firm acted independently.

The story which our model tells is, essentially, a simple one. Consider a trucking duopoly in which both firms overload some fraction of their trucks on a particular route. Firm A's overloading induces an increase in the regulatory effort of the government on that route. This reduces the expected profits of Firm B, the other duopolist. Firm A can be said to inflict a 'regulatory externality' on Firm B. At the same time, however, Firm B inflicts an analogous regulatory externality on Firm A. We demonstrate that, in general, there will be benefit to both firms from agreeing to some kind of mutual restraint on their respective violation rates. Such cooperative action on the part of the members of the industry is what we will call 'self-regulation'. The new (cooperative) equilibrium will not be characterised by full compliance, but will involve less frequent violation than in the non-cooperative equilibrium. For this reason the government regulatory agency should encourage this class of cartelisation by the industry.

The focus of this study is the problem faced by firms in the industry, and the regulatory framework is to some extent 'black boxed', though we will motivate our assumptions regarding the characteristics of that black box in several ways. In Section 2 we develop and solve a simple model of the loading problem faced by the individual truck operator. In Section 3 we consider the role of market structure on the level of compliance of the industry. It is shown that a less concentrated industry will, in the aggregate, violate the weight standards more often than a more concentrated one. In Section 4 the scope for, and practicability of, self-regulation is examined. It is argued that the gains from industry self-regulation may be considerable, though empirical determination of those gains is likely to be hampered by the problem of adequately defining the boundaries of the market (how far the industry is segmented, both geographically and by the characteristics of the vehicle population). Suggestive examples taken from the North American industry are used throughout.

(2) A Stylised Model of the Truck-Loading Decision.

Weight regulation in a particular jurisdiction is based upon a weight limit, denoted W, which must not be exceeded. Loading a truck is a discrete choice problem: comply or overweight. The revenues and costs associated with complying or overweighting are, respectively, R_c , C_c , R_o, C_o. It is assumed that $R_c<R_o$ and $C_c<C_o$, but $(R_c- C_c) < (R_o- C_o)$. This last condition ensures that, in the absence of any regulatory enforcement, the operator faces an incentive to overweight: if it is not satisfied the problem is trivial.

The regulatory agency conducts random inspection of trucks in transit, in which permanent weigh-stations or patrol cars equipped with mobile scales are used to verify the truck's compliance status. The probability that a given truck on a given journey will be inspected in this manner will be denoted μ, and if the truck is found to be in violation of the weight limit its owner faces a fine of monetary value. The proportion of the ith firm's trucks which are overweighted is denoted β_i. The enforcement effort of the regulator, on any particular route or in any particular jurisdiction, is assumed to be an decreasing function of the propensity of the truck population to comply. If the industry in that segment is comprised of n firms indexed by lower-case subscripts, with respective market shares denoted s_i, then $\mu=\mu(\sum ß_i.s_i)$ where it is supposed that $\mu'>0$, $\mu''<0$. Increased disobeyance of the regulations on any particular routeway (or in any particular market segment) induces, ceteris paribus, increases in the intensity with which that routeway is policed. Such a relationship can be generated by variations on a variety of 'familiar' models of enforcement strategy from the literatures on environmental protection, income tax collection and others. These will be explored more fully in Section 3. For now we wish to focus on the problem from the point of view of the industry and treat the function $\mu=\mu(\sum ß_i.s_i)$ descriptively.

The trucking firm chooses a mixed strategy, described by the parameter ß which dictates the probability that the operator will dispatch an overweighted truck. Through time ß will come to be the proportion of truck-journeys which involve excessive loads. The term Ω will be used to denote the average rate of violation in the truck population as a whole. The problem faced by the ith firm, then, is to specify a mixed strategy (ie a ß such that $0 \leq ß_i \leq 1$) to maximise expected profit P_i.

(1) $\text{Maximise}_{ß}\ P_i$

$$=\{[1-ß_i].[R_c-C_c] + ß_i.[R_o-C_o] - ß_i.\mu(\Sigma s_i.ß_i).f\}$$

The first and second terms represent expected net operating surplus from compliant and overweighted trucks respectively, whilst the third term is the expected penalties to be paid. Let $G=[R_o-C_o]-[R_c-C_c]>0$, the increase in gross operating profit (ie abstracting from regulatory penalty considerations) generated by overweighting a truck which would otherwise be in compliance. In the case of an interior solution and assuming Nash conjectures (ie that $(\partial ß_j/\partial ß_i)=0$ for all $i\neq j$, the firm takes the actions of its competitors as given), the first-order and second-order conditions (where the latter is assumed to be satisfied) associated with the ith firms problem are

(2) $(\partial P_i(ß_i^*)/\partial ß_i)$

$$= G - \mu(\Sigma s_i.ß_i^*).f - ß_i^*.\mu'(\Sigma s_i.ß_i^*).f.s_i = 0,$$

and

(3) $(\partial P_i^2(ß_i^*)/\partial ß_i^2)$

$$= - f.s_i.\{2.\mu'(\Sigma s_i.ß_i^*) + ß_i^*.s_i.\mu''(\Sigma s_i.ß_i^*)\} < 0$$

respectively, where the asterisk denotes the solution value.

The first-order condition indicates that the firm will go on increasing the proportion of its trucks that it overweights up to the point at which the expected increase in gross profit from so doing equals the increase in expected penalties. The increase in penalties has a direct and an indirect component. There is an increased likelihood that if one of that firms trucks is inspected (which will be the case with probability μ) it will turn out to be overweight and thus draw a fine (of size f). In addition, the simultaneous increase in the proportion of non-compliant trucks on the road will lead to an increase in the enforcement effort of the regulator and thus increase the risk of detection of all of that firm's (and, incidentally, every other firm's) overweight trucks. There is evidence that truck-operaters do trade-off the costs and benefits from breaching weight requirements in this way when travelling on any particular routeway. One study records that ". . . truckers often have substantial experience to aid them in calculating the probability of being apprehended. Based on this probability they can calculate the expected costs (in fines) of overweighting", (Paxson and Glickert (1982: 34)).

Assuming that the optimal mixed strategy for firm i is described by an interior solution for ß (that is, the firm dispatches some, but not all, of its trucks overweight) application of the implicit function theorem to the first-order condition (Equation 2) yields several comparative static results: $(\partial \beta_i^*/\partial R_c)<0$, $(\partial \beta_i^*/\partial R_o)>0$, $(\partial \beta_i^*/\partial C_c)>0$, $(\partial \beta_i^*/\partial C_o)<0$, $(\partial \beta_i^*/\partial f)<0$,

$(\partial \beta_i^*/\partial \mu)<0$ and $(\partial \beta_i^*/\partial \beta_j)<0$, all of which have straightforward interpretations and are strongly intuitive. A particular firm is more likely to violate when G (the gross gain associated with violation) is large, fines are small, inspection frequencies are low and/or its competitors are violating relatively infrequently.

(3) Market Share and Overweighting.

The vector of n first-order conditions (or, where appropriate, corner solutions), one for each firm in the market segment which we are examining, implicitly serves to define the compliance behaviour of each individual firm, and hence the environmental performance of the industry as a whole. It is of some interest to consider both the impact of a particular firm's market share upon that firm's respect for the weight regulations and, more generally, the impact of market structure upon the compliance of the industry in the aggregate. This is of particular interest in considering the possible impact of industry deregulation upon the integrity of the road network.

The inter-city road-haulage industry in Canada (as in most other developed economies) contains firms of many different sizes - from corporations with over a thousand trailers to independent operators with a small handful. Transport Canada data can be used to estimate four-firm concentration ratios for road-haulage services in the Canadian provinces (see Hirshorn

(1981: 51)). The total figures range from 15% (for Quebec and Ontario) to as high as 46 and 52% for Manitoba and Saskatchewan respectively. The licences granted to firms are, in most cases, highly specific with regards to the types of vehicles and the precise routes which the operator can utilise.[1] The market structure for freight-haulage on any given routeway can, therefore, be significantly more concentrated than this.[2] The hypothesis of this paper is that operators in a particular market segment (ie on a given routeway) are engaged in a non-cooperative game. The larger firms, in the cases mentioned in footnote 3 for example, will take account of the effect which their compliance decisions will have upon the regulatory enforcement effort and thus upon the profitability of their whole fleet of trucks. The operator of a single truck has comparatively little incentive to do this and will be more prone, according to our analysis, to overload.

In Sections (3.1) and (3.2) the impact on the firm's decision problem of an increase in its market share is analysed in three different contexts. In each case the market share of a particular

---

(1) See, for instance, Boucher (1981) for details of the limitations on the permits issued to hauliers on the principal Montreal to Quebec City route.
(2) Hirschorn (1981: 53) records the number of competitors on various different routeways in Canada. For one example amongst many, from Sault St. Marie to Sudbury there were six competing firms: four were large (with more than 600 trailers), one was medium (300 to 600) and one was small (less than 300). On the primary Toronto to Ottawa route there were nine: two large, four medium and three small.

haulage firm is taken as given; market structure is assumed not to be sensitive to changes in the régime of weight regulation in place.

(3.1) Market Share and Overweighting: General.

Consider the case of a symmetric, n-firm Nash equilibrium. In this case the market share of the representative firm is (1/n). Defining Ω to be the industry non-compliance rate (that is, $\Omega = \sum s_i.ß_i$) the first-order condition associated with an interior solution to the ith firm's truck-loading problem is, from Equation 2,

$$(4)\quad (\partial P_i(ß_i^*)/\partial ß_i) = 0$$
$$= G - \mu([ß_i^* + \sum ß_h]/n).f - ß_i^*.\mu'([ß_i^* + \sum ß_h]/n).f/n.$$

The assumption that the ith firm has Nash conjectures regarding the reactions of rival firms to changes in his actions means that $(\partial ß_h/\partial ß_i)=0$, such that $(\partial \Omega/\partial ß_i)=s_i=(1/n)$. It will be assumed, again, that the requisite sufficiency condition is satisfied. In equilibrium, symmetry requires that $ß_i=ß_h$, such that the subscripts which distinguish the firms can be dropped and the representative first-order condition becomes:

(5) $G - \mu(\beta^*).f - \beta^*.\mu'(\beta^*).f/n = 0.$

To examine the impact of market structure upon the propensity of the industry to overload its trucks, the implicit function theorem can be applied to this equation to show that,

$$(6)\ (\partial\beta^*/\partial n) = -(1/J).\{\beta^*.f.\mu'(\beta^*)/n^2\}$$

where J represents $\partial^2P(\beta_i)/\partial\beta_i^2$ in the representative firm's problem and, assuming that the pertinent second-order condition holds, is negative in the vicinity of β^*. The term in braces is unambiguously positive. Thus, $(\partial\beta^*/\partial n)>0$: an <u>increase</u> in the number of constituent firms in a symmetric, n-firm Nash equilibrium industry will be associated with an <u>increase</u> in the proportion of trucks which each firm (and, thence, the industry as a whole) dispatches in violation of the weight regulations.

There <u>may</u> exist, depending on the form of the regulatory reaction function, $\mu(\Omega)$, some N_{min} and N_{max} such that $\beta^*=0$ when $n<N_{min}$ and $\beta^*=1$ when $n>N_{max}$. That is, when the number of firms in the industry is sufficiently small there will be no overweighting in equilibrium (unless N_{min} turns out to be less than one, in which case this possibility is eliminated). In that case, when the number of firms in the industry is sufficiently large all trucks will be overweighted. Notice that N_{max} will be finite only if $G>\mu(1).f$, (ie if $\mu(1)$ is sufficiently small).

When one firm overweights one of its trucks it induces, by increasing the fraction the industry population which is overweighted, an increase in the intensity of regulation. By so doing it inflicts a 'regulatory externality' on its rivals. Increasing the number of firms in the industry increases the portion of costs due to the regulatory tightening which constitute external costs. As its share of the industry-wide regulatory costs of overweighting decreases (ie as the industry becomes less concentrated) the firm is inclined to internalise a smaller part of the regulatory costs which it imposes.

In the symmetric case the market share of each firm, s_i, equals (1/n). Thus, more overweighting is associated with firms with smaller market share. This result can be generalised to an asymmetric case, in which the industry is populated by firms of different sizes. The analytic result that propensity to overweight is a decreasing function of the size (market share) of the carrier conforms to the evidence that it is the small, independent operator who is prone to frequent violations. Walton and Yu (1983) provide empirical and anecdotal evidence of this in a study of road-haulage in Texas: "Common carriers. . .have low rates of violation. These data correspond well to comments rendered by

DPS personnel with respect to their observation that independent (small) truck operaters are the significant challenge to License and Weight officers" (Walton and Yu (1983: 28)).

(3.2) Compliance Performance and Enforcement Effort: Further Consideration.

Crucial to our model is the supposition that the enforcement efforts of the regulatory agency are 'reactive' in the sense that a change in the proportion of trucks not obeying the weight regulations will induce a change in the inspection rate. We have made the supplementary assumption that the relationship is positive and concave, such that the 'reaction function' of the regulator is upward sloping in $\{\mu,\Omega\}$-space.

It is not the aim of this paper to build in any kind of model of optimal regulatory behaviour (i.e., to derive the regulator's reaction function). Though the regulatory problem has not, to our knowledge, been subject to rigorous economic examination, a variety of familiar models could be used to derive a reaction function with the requisite characteristics. These include (i) simple 'crime-and-punishment' models in which the increased presence of non-compliers in the regulated population increases the productivity of the marginal dollar spent on inspection, thus causing more to be spent upon it; (ii) 'segmented market' models in which a national or provincial agency allocates some fixed budget across a set of market segments (ie different routeways) to maximise the number of violators caught or some other objective function; (iii) models of a budget-constrained agency in which penalty income is 'recycled' (either directly or indirectly) into the agency's operating budget; (iv) 'transport planning' models, in which a central planning department (such as a Provincial Highway Commission) with overall control of maintaining the integrity of the road network finds it cost-effecive tc respond to increased non-compliance rates by redirecting funds from its maintenance budget (ie damage repair) to its enforcement budget (ie damage prevention); (v) 'behavioural models' in which for (often ill-defined) bureaucratic reasons some desired compliance rate has been targetted by administrative or legislative mandate; and (vi) various 'political economy' models in which increased 'crime' rates and diminishing road quality standards increase public and political awareness of the problem, and thus the efficacity of the various automobile, environmental and other interested pressure groups.

This list of 'stories' which could be told to generate a functional form for $\mu(.)$ of the type assumed in our analysis is suggestive rather than exhaustive. Furthermore, they are not mutually exclusive; two or more effects could work at the same time. It should be straightforward, conceptually at least, to use time-series methods to estimate the relationship. In the next two sub-sections we reasses the impact of market structure upon respect for weight regulations in the case of a budget constrained enforcement agency. This corresponds to case (iii) in the last paragraph, and is the type of arrangement found in some jurisdictions in France.

(3.2.1) A Budget-Constrained Enforcement Agency: Opening the Regulatory 'Black Box'.

Suppose that the budget that the enforcement agency receives from the central government is equal to b units per truck-trip made in the jurisdiction. The amount is, presumably, some fraction of the constant unit cost of executing an inspection, denoted k. Thus $k>b$ except in the trivial case of a comprehensive inspection programme. In addition to this, the agency receives the income from penalties levied. The operating budget of the enforcement agency per truck-trip is denoted B, and can be described by

$$(7)\ B = b + \Omega.\mu(\Omega).f,$$

the sum of its income from government grant and from penalty income.

Again we restrict our attention to the case of a symmetric, n-firm industry. The agency spends all of its revenues on instigating random inspections and there is no fixed cost associated with the inspection programme. Choosing μ, the inspection rate, subject to this budget constraint (which is assumed to bind) and the reaction function of the industry means that $\mu(\Omega)$ is implicitly defined. Suppose, momentarily, that N trips, of which M are subject to inspection. Then balancing the regulatory agency's budget implies that

(8) $k.M = (b + \Omega.\mu(\Omega).f).N$

Rearranging terms and recognising that $\mu=(M/N)$ the equilibrium probability of a truck's being inspected on a particular journey is

(9) $\mu(\Omega) = (b/[k - \Omega.f])$

This, then, constitutes the reaction function of the budget-constrained enforcement agency to increased industry non-compliance. The agency can be thought of as an automaton which takes the revenue from penalties levied and 'recycles' it into a verification programme, where the intensity of that programme is constrained by the agency's 'grant plus penalty' budget. The reaction function thus generated has the requisite characteristics (in particular its first and second derivatives are positive and negative respectively). Notice that, because the government grant component of the inspection budget is predetermined and positive, $\mu(0)$ is positive in this case.

Without further elaboration, it can be shown that the first-order condition of the representative ith firm in the symmetric n-firm industry reduces to:

$$(10) \quad (\partial P_i(ß_i^*)/\partial ß_i) = 0$$
$$= G - f.(b/[k - \Omega.f]) - (f.ß_i^*/n).(f.b/[k - \Omega.f]^2)$$

It is straightforward to demonstrate that the second-order condition required for a maximum is satisfied everywhere.

If G (the gross gains from violating the standard) is sufficiently small vis-à-vis the expected penalty at the lowest rate of regulatory inspection then all firms will always comply, regardless of market structure. In particular, if

$$(11) \quad G - \mu(0).f < 0$$

or, substituting for $\mu(0)$,

$$(12) \quad G - b.f/k < 0,$$

then $ß^*$ (and thence Ω) equals zero for all (positive) values of n. This would imply that no operator would ever overload a truck. This condition trivialises the problem and, for empirical as well as analytic interest, it will be assumed not to hold; there is some non-zero equilibrium rate of violation. In the case of an interior solution to the representative firm's problem (and recognising that symmetry is required of the solution, ie $ß_i=ß_h=\Omega$) it is true that

$$(13)\quad \partial\Omega/\partial n = -(1/J).f^2.\Omega.b \Big/ \left\{ n^2.[k - \Omega.f]^2 \right\} > 0,$$

where J has the usual interpretation and is negative. The sign of the derivative in Equation 13 is, then, determined by the sign of the expression in square brackets, which is unambiguously positive. Thus $\partial ß^*/\partial n$ is positive; an increase in the number of firms in the industry worsens its compliance performance and induces greater inspection efforts. Sufficient entry into the industry may cause non-compliance with the regulations to become chronic. Analytically, the Kuhn-Tucker conditions imply that $ß^*=1$ when $n>N_{max}$ where N_{max} is implicitly defined by

$$(14)\quad G - \mu'(1).f/N_{max} - \mu(1).f = 0.$$

After substitution,

$$(15)\quad N_{max} = b.f^2 \Big/ \left\{ [k-f].[G.(k-f)-b.f] \right\}$$

The results derived here are qualitatively similar to those derived with a general. descriptive functional form in Section 2. In this case, however, a 'story' has been told to rationalise the regulatory reaction function employed descriptively in earlier analysis.

(4) The Benefits From Industry Self-Regulation.

In this section we show that there exist, in general, gains to the industry from self-regulation. In addition, we present a diagram of the equilibrium for the case of a duopoly, and discuss the implementability of a self-regulatory régime.

Recall that P_i is the expected profit of the representative firm. Consider a small chage in Ω, generated by each firm's simultaneously changing its choice of ß by the same arbitrarily small amount. In the vicinity of the (non-cooperative) Nash equilibrium,

$$(16)\quad dP_i(\Omega^*)/d\Omega = (\partial P_i(\Omega^*)/\partial\Omega) + (\partial P_i(\Omega^*)/\partial\mu).\mu_\Omega(\Omega^*)$$

The full impact of the change upon the expected profits of the ith firm is the sum of two components. The first is the 'direct effect', holding the level of regulatory enforcement constant. This term is positive, except in the case of n=1 (a monopolised industry) in which case the envelope theorem dictates that it should equal zero. The second term is the effect on expected profitability of the change in regulatory stringency induced by the change in the proportion of violators on the road. This effect will be unambiguously negative, as marked. It can be shown that the indirect effect always outweighs the direct effect by re-expressing Equation 16 (the substitution for the 'direct effect' can be made from the firm's first-order condition, Equation 2) as

$$(17) \quad dP_i(\Omega^*)/d\Omega = (1-s_i).\Omega.f.\mu_\Omega(\Omega^*) - \Omega.f.\mu_\Omega(\Omega^*)$$

$$= - s_i.\Omega.f.\mu_\Omega(\Omega^*) < 0$$

or, in the n-firm symmetric Nash case considered throughout this paper,

$$(18) \quad dP_i(\Omega^*)/d\Omega = -(1/n).\beta^*.f.\mu'(\beta^*) < 0$$

The implication is that an agreement by all firms in the industry to reduce (each by an equal amount) the frequency with which they overweight their trucks would increase the expected profit of <u>all</u> firms. The result is a variation on the well known Pareto inferiority of Nash equilibria. In the non-cooperative context examined, no firm has an incentive to take account of the 'regulatory externality' which it imposes upon the other firms in the industry. Thus, every firm overweights more frequently than it would if it were to take account of (or be liable for) the external costs it generates. The result is that the industry ends up in a 'frequent violation, high enforcement' equilibrium.

(4.1) Self-Regulatory Agreements: The Duopoly Case.

It is easiest to discuss the potential gains from self-regulation for the case of an industry comprised of two (symmetric) Nash competitors. In $\{ß_1,ß_2\}$-space it is possible to characterise a reaction function for each firm, of the form $ß^*_1=f(ß^*_2)$ and $ß^*_2=f(ß^*_1)$. It is straightforward to demonstrate that each of these will be downward sloping (and truncated at $ß=1$). There is, of course, no reason to suppose that the reaction functions will be linear - further characterisation requires that functional forms be imposed.

A core of $\{ß_1,ß_2\}$ pairs, which ensure greater expected profits to <u>both</u> of the firms than the Nash equilibrium, can be identified. Within the constraints of our analysis (ie that we only consider the welfare of the members of the industry) such a solution Pareto dominates the Nash equilibrium. There is, however, a third group whose interests have been suppressed in our analysis, namely the general public - in whose interest regulations are supposedly framed in the first place. It is reasonable to suppose, however, that this third group also prefer the new solution over the Nash equilibrium, since it is associated with a lower rate of violation (and thus road damage) and/or a lower rate of enforcement expenditure. The points in the 'core' can, then, truly be said to Pareto dominate the Nash equilibrium.

(4.2) The Implementation and Sustainability of a Self-Regulatory Accord.

It is not enough to identify potential gains from cooperative action, in any given context. It is also necessary to examine the extent to which those gains can be reaped. The practical problems associated with implementation of the requisite cooperation may turn out to be formidable.

The procedure whereby the firms in the industry attempt to implement a point in the 'core' (rather than reverting to the non-cooperative Nash solution) is what we will term 'self regulation'. Some convention must be arrived at whereby all operators in the particular geographic market segment agree to comply with the weight regulations more frequently than would be the case in the absence of self-regulation. By how much each firm would be required to raise its own, particular compliance rate cannot be determined uniquely. This reflects the fact that there are a multiplicity of possible solutions involving self-regulation. Which of the possible points in the core will be 'chosen' by the industry to constitute the solution will depend on, amongst other things, the institutional structure of the industry and the internal framework set-up to facilitate inter-firm bargaining over the terms of self-regulation.

The self-regulated outcome will be associated with less frequent violation and lower enforcement intensity than the non-cooperative Nash solution. In effect, the self regulatory programme has displaced some part of the external regulatory programme and, in this sense, self regulation by the industry can be considered a legitimate (partial) substitute for administrative regulation. The potential Paretian improvements generated by replacing some part of the latter with the former have rested, so far, on the implicit asssumption that enforcement within in the industry would be costless, whereas enforcement from outside is costly. The case can clearly be made, however, whenever internal enforcement is less costly than external enforcement (though not necessarily costless).

The central problem associated with implementing an industry equilibrium with self-regulation is the familiar issue of sustainability. (See Shubik (1975) for a classic textbook treatment, also various chapters of Schmalensee and Willig (1989).) From any point in the core identified above, every firm stands to increase its expected profits by unilaterally defecting from that point. More generally, this is true for any point within the area bounded by the reaction functions and the axes in $\{\beta_1, \beta_2\}$-space, with the exception of that associated with the Nash equilibrium. Unless a more cogent story is told as to why the representative firm would not 'cheat' on any industry agreement it is unclear why any solution point other than the Nash equilibrium should be seriously considered. Even if the firms in the industry did agree to some mutual reduction in violation rates it would not be sustainable - we would soon expect to find operators breaking the terms of the agreement and the self-regulatory system breaking down.

It is, then, necessary that we 'tack on' some additional tool whereby the firms in the industry will be able to sustain mutually beneficial cooperation. Our recommendation is the establishment of 'cartel fines'. Whenever the truck of a particular operator is fined for overweighting (news of such prosecutions are in the public domain) that operator would, as well as the f dollars paid to the formal regulatory agency, pay an additional 'cartel fine' equal to f_c dollars to the internal regulator. The funds thus generated would be paid to the the local truckers association or other nominated self-regulatory body.

The cartel fine which would implement the optimal (ie joint profit maximising) solution would be equal to

$$(19)\quad f_c = [(1-s_i)/s_i].f.\mu'(\Omega^*).\Sigma_{-i}\beta_i\ ,$$

Where the summation in the expression is done across all of the firms in the industry except for the ith. In the symmetric case this becomes

$$(20)\quad f_c = (n-1).f.\mu'(\beta^*).\beta^*$$

The cartel fine, f_c , constitutes an optimal 'Pigovian' fine and serves to internalise the regulatory externality which each firm would otherwise be visiting upon its competitors. Notice, from Equation 20, that the appropriate f_c is everywhere positive and increasing in n. As the representative firm's market share increases the magnitude of f_c diminishes - capturing the notion that a large firm generates a smaller externality. When n=1 (ie s_i=1) the industry is monopolised and the appropriate cartel fine would be zero - the monopolist generates no 'regulatory externality' because there is no other firm in the industry upon which such an externality could be incident.

Though we have restricted our attention to symetric industry configurations, the results can be shown to be generalisable in an intuitive way to non-symetric cases. In an industry populated by firms of different sizes the biggest cartel fine would, by extension of the results above, be payable by the small operator with a single trailer. Smaller fines would be paid by the larger firms. It is questionable that a 'regressive' system of this type would sit well with advocates of deregulation. Seeming to penalise the small firm disproportionately harshly, it might be interpreted as an attempt by large industry members to increase their market dominance. The best we might hope for would be flat-rate cartel fines (the distinction is, of course, trivial in the symmetric market structure case which was the focus of the earlier analysis).

Whether this type of self-regulation proved workable would depend, to a large extent, on institutional considerations. Implementability would vary among jurisdictions. There is an implicit assumption that some channel exists whereby each firm can be coerced into paying the requisite cartel fine - that by failing to do so it would be expelled from the Association and that expulsion losing the other benefits from membership. These other benefits may, in some cases, be too small to permit expulsion to constitute a serious enough sanction to , and this could hinder workability. One way around this might be to require each operator to deposit some fee with its association, which would be forfeited if the operator failed to honour the terms of the self-regulatory agreement.

Economic theory predicts that implementation of 'cooperative' equilibria in a non-cooperative setting is likely to be particularly difficult when there are many players involved. It is also likely to be difficult to sustain cooperative action when the jurisdiction is host to a large number of 'transient' trucks which may use the routeways but are not party to the local self-regulatory accord. In examining the feasibility of a self-regulatory régime, then, it is important that the market be adequately delineated.

At a provincial or regional level the appropriate institutions for administering a self-regulatory framework (based on cartel fines or some other mechanism) would be the relevant trucking Associations.[3]

One province in which cooperative action be the trucking industry could be expected to be practicable is Newfoundland. This is particularly true following the recent evolution of the market's structure there. Previously, the industry was highly fragmented with many small firms providing service which was not always efficient. With the emergence of of larger firms and a considerable reduction in the number of one-man carriers, a group of twelve carriers has now captured almost all of the trucking market in Newfoundland (the Sullivan Commission noted the rapid rate of small firm consolidations and the subsequent rise to predominance of "10 to 12 large carriers"). This relatively small 'club' of operators, along with the lack of any 'through' traffic (for obvious geographical reasons) of trucks belonging to operators outside of the jurisdiction of the local industry federation makes it particularly likely that a mutually advantageous self-regulatory accord (based, perhaps, on 'cartel fines') could be implemented in that province.

---

(3) In Canada there are seven of these (the Manitoba, Atlantic Provinces, Ontario and Saskatchewan Trucking Associations, the British Columbia and Alberta Motor Transport Associations and L' Association du Cammionage du Quebec), each affiliated with the Canadian Trucking Association. The position of smaller National organisations which represent smaller sections of the truck population (eg the Canadian Association of Movers or the Private Motor Truck Council of Canada) in a self-regulatory agreement is less clear. Comprehensivity, or near comprehensivity, would be vital to the success of such a scheme.

It may be, however, that the natural level of aggregation for such self-regulatory agreements is not at the provincial or regional level. The detailed restrictions on routes to be travelled on a given operators licence means that the industry can be treated as a highly segmented one, and it may be that agreements could be drawn-up at the segmental level. Thus, it is conceivable that the six competing firms on the Sault St. Marie to Sudbury route, mentioned earlier, could arrive at some agreement regarding overweighting on that route (ie in that market segment). In response to these cuts in violation frequency the government regulator would (rationally) cut the enforcement effort which he exerts on that routeway and, as our analysis has demonstrated, all of the six firms could expect to gain.

Cooperative action in any given market segment is also likely to be easier to sustain when the gains to the various parties from cooperation are well-defined. It would, we contend, be in the interests of the government agency to facilitate self-regulatory agreement by making the linkage between observed violation rates, on any given routeway, and future rates of inspection on that routeway more explicit. This would involve commitment to an announced enforcement policy which would serve, analytically, to render the μ-function common knowledge and, in so doing, make the gains to the industry from self-regulation more transparent.

(5) Conclusions.

In this paper we have shown that, in general, there are gains to be had by the trucking industry to (at least partial) self-regulation of truck weight. These arise because of the existence of <u>regulatory</u> externalities.

The analysis is restricted to the case in which self-regulation is desireable, from the point of view of the industry, because external supervision is inadequate. In the model this inadequacy is generated by the low level of fines for those caught violating weight limits (as we have noted fines oserevd in the US and Canda are, indeed, extremely small, vis-à-vis the level which would give truckers the incentive to comply with regulations). These were assumed to be fixed according to exogenous criteria (the assessment of courts of the graviy of the crime committed). If the level of penalties were optimally set, in conjunction witinspection probabilities, by the external industry watchdog then, in all probability, the need and scope for industry seflf-regulation would disappear.

Local feasibility of a self-regulatory programme is likely to depend upon local institutional factors. Case-study analysis of implementability by region (or even by well defined routeway) is a possible next step. The main contribution of this paper has been to use economic theory to demonstrate that there are, in all but the most trivial cases, gains to be reaped by the industry from self-imposed regulatory programmes. These can be expected, in general, to take the form of 'codes of practice' and should be backed-up by credible monetary sanctions. We have also argued that such programmes will be socially beneficial and, as such, should be encouraged by government regulatory bureaux.

Conclusions.

The theme of the papers in this thesis has been the lack of 'easy' answers in the regulation of environmental externalities (and, indeed, externalities more generally).

Chapter 1 explored in some detail the complications which incomplete compliance with a standard can have on the analysis of regulatory reform. In Chapter 2 we examined, in an explicitly dynamic framework, the complications for environmental standard-setting associated with the endogeneity of the supply of new environmental technologies. In both cases the frailty of some parts of conventional wisdom in the domain of environmental regulation was demonstrated. In Chapter 3 we uncovered a class of Pareto improvements to be reaped from self-regulation on the part of an already regulated industry. The feasibility of executing those improvements was argued, however, to be hampered by various implementation problems.

More detailed conclusions drawn from the analysis were presented at the end of the appropriate chapter. They all point to the complexity which must characterise adequate discussion of environmental regulatory issues.

Chapter 1: Bibliography.

Allingham, M. G. and Sandmo, A. 1972 "Income Tax Evasion: A Theoretical Analysis". Journal of Public Economics (1): 323-338.

Alm, J. 1988. "Compliance Costs and the Tax-Avoidance, Tax-Evasion Decision. Public Finance Quarterly (16): 31-66.

Baiman, S. and Demski, J. S. 1980. "Economically Optimal Performance Evaluation and Control Systems". Journal of Accounting Research (18, Supplement): 184-220.

Baron, D. P. and Besanko, D. "Regulation, Asymmetric Information and Auditing". Bell Journal of Economics (15): 447-470.

Becker, G. 1968. "Crime and Punishment: An Economic Analysis." Journal of Political Economy (76): 169-217.

Besanko, D. 1987. "Performance Versus Design Standards in the Regulation of Pollution". Journal of Public Economics (34): 19-44.

Border, K. and Sobel, J 1987. "Samurai Accountant: A Theory of Auditing and Plunder". Review of Economic Studies (54): 525-540.

Clotfelter, C. (1984). "Tax Evasion and Tax Rates: An Analysis of Individual Returns". Review of Economics and Statistics. (65): 363-373.

Craswell, R. and Calfee, J. E. 1984. "Some Effects of Uncertainty on Compliance with Legal Standards". Virginia Law Review (70): 965-999.

Craswell, R. and Calfee, J. E. 1986. "Deterrence and Uncertain Legal Standards". Journal of Law, Economics and Organisation (2): 279-303.

Downing, P. B. and Hanf, K. 1983. "International Comparisons in Implementing Pollution Laws". Kluwer Nijhoff Publishing, Boston.

Downing, P. B. and Watson, W. D. 1974. "The Economics of Enforcing Air Pollution Controls". Journal of Environmental Economics and Management (1): 219-236.

Erekson, O. H. and Sullivan, D. 1988. "A Cross-Section Analysis of IRS Auditing". National Tax Journal (41): 175-189.

Graetz, M. J. and Wilde, L. L. 1985. "The Economics of Tax Compliance: Fact and Fantasy". National Tax Journal (38): 355-363.

Greenberg, J. 1984. "Avoiding Tax Avoidance: A Game-Theoretic Approach". Journal of Economic Theory (32): 1-13.

Harford, J. D. 1978. "Firm Behaviour Under Imperfectly Enforceable Pollution Standards and Taxes". Journal of Environmental Economics and Management (5): 26-43.

Harford, J. D. 1987. "Self-Reporting of Pollution and the Firms Behavior Under Imperfectly Enforceable Regulations". Journal of Environmental Economics and Management (14): 293-303.

Harford, J. D. and Harrington, W. 1991. "A Reconsideration of Enforcement Leverage When Penalties are Restricted". Journal of Public Economics (45): 391-395.

Harrington, W. 1988. "Enforcement Leverage When Penalties are Restricted". Journal of Public Economics (37): 29-53.

Jones, C. A. 1989. "Standard-Setting with Incomplete Enforcement Revisited". Journal of Policy Analysis and Management (8): 72-87.

Kambhu, J. 1989. "Regulatory Standards, Non-Compliance and Enforcement". Journal of Regulatory Economics (1): 103-114.

Kanodia, C. S. 1985. "Stochastic Monitoring and Moral Hazard". Journal of Accounting Research (23): 175-192.

Keeler, A. G. 1991. "Noncompliant Firms in Transferable Discharge Permit Markets: Some Extensions". Journal of Environmental Economics and Management (21):180-189.

Landsberger, M. and Meilijson, I. 1982. "Incentive-Generating State-Dependent Penalty System". Journal of Public Economics (19): 333-352.

Langbein, L. and Kerwin, C. M. 1985. "Implementation, Negotiation and Compliance in Environmental Safety Regulation". Journal of Politics (47): 854-880.

Lee, D. W. 1984. "Monitoring and Budget Maximisation in the Control of Pollution". Economic Inquiry (8): 565-575.

Leung, S. F. 1991. "How to Make the Fine Fit the Corporate Crime?" Journal of Public Economics (45): 243-256.

Linder, S. H. and McBride M. E. 1984. "Enforcement Costs and Regulatory Reform: The Agency and Firm Response". Journal of Environmental Economics and Management (11): 327-346.

Magat, W. A., Krupnick, A. J. and Harrington, W. 1986. "Rules in the Making: A Statistical Analysis of Regulatory Agency Behaviour". Resources for the Future, Washington, DC.

Magat, W. A. and Viscusi, W. K. 1990. "Effectiveness of the EPA's Regulatory Enforcement: The Case of Industrial Effluents Standards". Journal of Law and Economics (33): 331-360.

Malik, A. S. 1990. "Markets for Pollution Control When Firms Are Non-Compliant". Journal of Environmental Economics and Management (18): 97-106.

Mintz, J. A. 1988. "Agencies, Congress and Regulatory Enforcement: A Review of EPA's Hazardous Waste Enforcement Effort, 1970-1987". Environmental Law (18): 683-759.

Mookherjee, D. and P'ng, I. P. L. 1990. "Enforcement Costs and the Optimal Progressivity of Income taxes". Journal of Law, Economics and Organisation (6): 411-431.

Niskanen, W. A. 1972. "Bureaucracy and Representative Government". Aldine-Atherton Press, Chicago.

Niskanen, W. A. 1975. "Bureaucrats and Politicians". Journal of Law and Economics (17): 617-637.

Polinsky, A. M. and Shavell, S. 1979. "The Optimal Tradeoff Between the Probability and Magnitude of fines". American Economic Review (69): 880-891.

Portney, P.R. (ed.) 1990. "Public Policies for Environmental Protection". Resources for the Future, Washington DC.

Reinganum, J. F. and Wilde, L. L. 1985. "Income Tax Compliance in A Principal-Agent Framework". Journal of Public Economics (26): 1-18.

Reinganum, J. F. and Wilde, L. L. 1986. "Equilibrium Verification and Reporting Policies in a Model of Tax Compliance". International Economic Review (27):739-760.

Roth, J. A., Scholz, J. T. and Witte, A. D. 1989. "Taxpayer Compliance", (in two volumes), University of Pennsylvania Press, Philadelphia.

Russell, C. S. 1990. "Monitoring and Enforcement", Chapter 7 in P.R. Portney (ed.), "Public Policies for Environmental Protection". Resources for the Future, Washington DC.

Scotchmer, S. and Slemrod, J. 1989. "Randomness in Tax Enforcement". Journal of Public Economics (38): 17-32.

Segerson, K. 1988. "Uncertainty and Incentives for Nonpoint Pollution Control". Journal of Environmental Economics and Management (15): 87-98.

Snavely, K. 1990. "Governmental Policies to Reduce Tax Evasion: Coerced Behavior vs Services and Values Development". Policy Sciences (23): 57-72.

Stigler, G. 1970. "The Optimal Enforcement of Laws". Journal of Political Economy (78): 526-536.

Storey D. J. and McCabe, P. J. 1980. "The Criminal Waste Discharger". Scottish Journal of Political Economy (27): 30-40.

Strock, J. M. 1990. "Protecting the Environment: EPA After 20 Years". Trial (August): 22-27.

Townsend, R. M. 1979. "Optimal Contracts and Competetive Markets With Costly State Verification". Journal of Economic Theory (21): 265-293.

Tsebelis, G. 1991. "The Effect of Fines on Regulated Industries; Game Theory vs Decision Theory". Journal of Theoretical Politics (3): 81-101.

Viscusi, W. K. and Zeckhauser, R. 1979. "Optimal Standards with Incomplete Enforcement". Public Policy (27): 437-456.

Willick, D. H. and Windle T. J. 1973. "Rule Enforcement by the Los Angeles County Air Pollution Control District". Ecology Law Quarterly (3): 507-534.

Xepadeas, A. P. 1991 "Environmental Policy Under Imperfect Information; Incentives and Moral Hazard". Journal of Environmental Economics and Management (20): 113-126.

Chapter 2: Bibliography.

Ashford, N. A. and Heaton, G. R. Jr. 1983. "Regulation and Technological Innovation in the Chemical Industry". Law and Contemporary Problems (46), 109-155.

Baumgardner, J. R. 1991. "The Interaction Between Forms of Insurance Contract and Types of Technical Change in Medical Care", Rand Journal of Economics (22), 36-53.

Binswanger, H. and and Ruttan, V. 1978. "Induced Innovation". Johns Hopkins Press. Baltimore.

Cabral, L. M. B. and Riordan, M. H. 1989. "Incentives for Cost Reduction Under Price Cap Regulation", Journal of Regulatory Economics (1), 93-102.

Clemenz, G. 1991. "Optimal Price Cap Regulation", Journal of Industrial Economics (39), 391-408.

Cohen, W. M. and Levinthal, D. A. 1989. "Innovation and Learning: The Two Faces of R&D", Economic Journal (99), 569-596.

Downing, P. B. and White, L. J. 1986. "Innovation in Pollution Control", Journal of Environmental Economics and Management (13), 18-29.

Foray, D. 1991. "Repères pour une économie des organisations de recherche". Revue d'Economie Politique (September), 783-808.

________ 1991. "The Secrets of the Industry are in the Air", Research Policy (4).

Harris, G. G. and Pepper, R. B. 1986. "Application of ALARA to the Operation of Nuclear Power Stations", in Lakey and Lewins (eds.) 1987, 62-75.

Kaser, G. 1989. "Acceptable Nuclear Risk; Some Examples From Europe". European University Institute, Working Paper No. 89/409, San Dominico, Italy.

Kneese, A. V. and Schulze, C. L. 1978. "Pollution, Prices and Public Policy", Brookings Intitute, Washington DC.

Kydland, F. and Prescott, E. C. 1977. "Rules Rather Than Discretion: The Inconsistency of Optimal Plans", Journal of Political Economy (84), 473-491.

Lakey, J. R. A. and Lewins, J. D. (eds.) 1987. "ALARA: Principles, Practice and Consequences", Proceedings of the symposim on 'ALARA - Quantitative Optimisation Techniques for Radiation Protection in the Nuclear Industry' held by the Institution of Nuclear Engineers, London, 18 September 1986. Adam Hilger, Bristol, England.

McHugh, R. 1985. "The Potential for Private Cost-Increasing Technological Innovation Under a Tax-Based, Economic Incentive Pollution Control Policy", Land Economy (61), 58-64.

Magat, W. 1978. "Pollution Control and Technological Advance: A Model of the Firm", Journal of Environmental Economics and Management (5), 1-25.

________ 1979. "The Effects of Environmental Regulation on Innovation", Law and Contempory Problems (43), 4-25.

Mendlesohn, R. 1984. "Endogeneous Technical Change and Environmental Regulation", Journal of Environmental Economics and Management (11), 202-207.

____________ 1986. "Regulating Homogeneous Emissions", Journal of Environmental Economics and Management (13), 301-312.

Milliman, S. R. and Prince, R. 1989. "Firm Incentives to Promote Technological Change in Pollution Control", Journal of Environmental Economics and Management (17), 247-265.

Mummery, P. W. and Anderson, A. R. 1986. "ALARA In Practice", in Lakey and Lewins (eds.) (1987), 104-114.

Nelson, R. R. 1959. "The Simple Economics of Basic Research", Journal of Political Economy (67), 297-306.

__________ 1961. "Uncertainty, Learning and the Economics of Parallel Research and Development Efforts", Review of Economics and Statistics (XLII), 351-364.

__________ 1982. "The Role of Knowledge in R&D Efficiency", Quarterly Journal of Economics (August), 453-470.

Note, 1979. "Environmental Rule-Making in Action - Setting a Standard for Coal Emissions", Resources (January), 15-16.

Orr, L. 1976. "Incentives for Innovation as the Basis of Effluent Charge Strategy", American Economic Review (56), 441-447

Spence, A. M. 1980. "The Learning Curve and Competition", Bell Journal of Economics (11), 49-70.

Webb, G. A. M. 1986. "The Development of ALARA -Philosophy and Practice", in Lakey and Lewins (eds) (1987), 3-23.

Weitzman, M. 1974. "Prices versus Quantities", Review of Economic Studies (41), 477-498.

Wenders, J. T. 1975. " Methods of Pollution Control and the Rate of Change of Pollution Control Technology", Water Resources Research (11), 393-396.

Wright, J. K. 1986. "Introduction to Session 1", in Lakey and Lewins (eds) (1987), 2.

Yao, D. A. 1988. "Strategic Responses to Automobile Emissions Control: A Game Theoretic Analysis", Journal of Environmental Economics and Management (15), 419-438.

Zhou, L. Y. and A. H. Rubenstein, 1986. "Imbedded Technology Capability (ITC) and the Management of Science and Technology in China: A Research Note", Research Policy (15), 49-51.

Chapter 3: Bibliography.

Barron, R. T., 1985. "Analysis of Pavement Damage Attributable to Overweight Trucks in New Jersey", Transportation Research Record 1038.

Boucher, M., 1982. "The Quebec Trucking Industry Regulation: A New Examination", Centre for the Study of Regulated Industries, Working Paper #1982-25, Montreal, Quebec.

Cubukgil, A., S. Flott and F. P. Nix, 1984. "A Conceptual Framework for Monitoring Transborder Trucking", University of Toronto/York University Joint Program in Transportation, Working Paper #14.

Euritt, M. A., 1985. "Overweight Truck Operation in Texas: Developing Effective Fines" Professional Report, Faculty of Graduate School, University of Texas, Austin.

Hilderbrand, M. D., B. E. Prentice and I. Lipnowski, 1991. "Enforcement of Highway Weight Regulations: A Game Theoretic Model", Transportation Research Record.

Hirshorn, R., 1981. "Trucking Regulation in Canada: Review of the Issues", Economic Council of Canada, Working Paper #26, Ottawa, Ontario.

Kaplan, H. 1989. "Policy and Rationality: The Regulation of Canadian Trucking Regulation", University of Toronto Press, Toronto, Ontario.

Organisation for Economic Cooperation and Development, 1983. "Impacts of Heavy Freight Vehicles", Paris, France.

Paxson, D. S. and J. P. Glickert, 1982. "Value of Overweighting to Intercity Truckers", Transportation Research Record 1038.

Schmalensee, R. and R. Willig (Eds.), 1989. "Handbook of Inductrial Organisation: Volume 1". Elsevier Science Publishers B.V., North Holland.

Shubik, M., 1959. "Strategy and Market Structure". Wiley, New York.

Tirole, J. 1989. "The New Industrial Organisation", MIT Press, Cambridge, MA.

Walton, C. M. and C. Yu, 1983. "Truck Size and Weight Enforcement: A Case Study", Transportation Research Record 920.

Wyatt, J. J. and M.U. Hassan, 1984. "Overweight Trucks: Incidence, Costs to the Highway System, Enforcement and Compliance", Transportation Forum 1(2).

Wyatt, J. J. and M.U. Hassan, 1985. " Some Tentative Findings About the Effect of Enforcement on Compliance with Weight Regulations", Transportation Forum 2(3).

U. S. General Accounting Office, 1979. "Excessive Truck Weight: An Expensive Burden We Can No Longer Support", Report #CED-79-94, Washington DC.