Center for Research on Economic and Social Theory
Research Seminar in Quantitative Economics

Discussion Paper

DEPARTMENT OF ECONOMICS
University of Michigan
Ann Arbor, Michigan 48109

[redacted]

Optimal Regulation of Research and Development Under Imperfect Information

David Sappington
University of Michigan

C-45

May 1981

Revised December 1981

Abstract

The optimal regulatory strategy to promote research and development aimed at cost reduction is derived in an environment where the firm's information about the technology of cost reduction, although initially imperfect, is better than that of the regulator. The manner in which the optimal regulatory strategy varies with changes in the informational environment is also described.

The author is indebted to Bengt Holmstrom, Wesley Magat, Hal Varian, Robert Willig, two anonymous referees, and the Editorial Board of The Bell Journal for their helpful comments and suggestions. These persons are, of course, in no way responsible for any errors that may remain.

1. Introduction

Of primary concern to many regulators is the extent to which the firms that they regulate are perceived to be innovative and actively engaged in reducing their costs of operation. Of course, regulated firms will have an interest in reducing costs only if the proper incentives are created for them to do so. It is the intent of this research to specify the properties of an optimal regulatory policy designed to foster cost reduction by a regulated firm under conditions of imperfect and asymmetric information.[1] The informational environment considered here is one in which neither the firm nor the regulator initially know with certainty the level of cost reduction that will result from any specified amount of research and development effort on the part of the firm. The firm, however, does always possess better information (in the sense of Blackwell [1951]) than that of the regulator about the likely consequences of its efforts designed to reduce costs. It is only after a research strategy has been chosen (in response to the incentive scheme offered by the regulator) that the firm discovers the most efficient manner to realize any specified alteration of the cost function.

In order to focus on the effects of imperfect and asymmetric information about the technology of cost reduction, other information is assumed to be perfect and symmetric. In particular, both the firm and the regulator are assumed to know the cost and demand functions that face the firm at any moment in time. Thus, any modification of the cost function effected by the firm can be observed by the regulator and translated into a single number that corresponds to the maximum possible gain in consumers' surplus that can be realized as a result of the observed change in costs.[2]

Furthermore, the analysis here is essentially static so as to rule out the possibility of strategic dynamic interaction between regulator and firm (of the type considered by Sappington [1980]). Thus, there are no gains to the firm, for example, from deliberately expending too little effort to reduce costs today in order to gain greater benefit under subsequent incentive schemes initiated by the regulator in future years.

The analysis is also made less complicated, but perhaps more realistic, by restricting the principal to linear incentive schemes. In other words, payments to the firm for realized cost reductions must be comprised of a fixed payment independent of the value of such reductions plus a prespecified fraction of this value. In practice, regulatory rules are often expressed as simple formulae rather than complicated functional relationships.[3]

The analysis will proceed as follows. In Section 2, the model is described more completely and expressed formally. The properties of the optimal linear incentive scheme are then derived in Section 3. It is shown that the firm will be permitted to choose among compensation schedules, all but one of which will result in a final outcome that is ex post Pareto inefficient (even though the regulator could always ensure an efficient outcome if he chose to do so). An intuitive explanation of the findings is offered in Section 4 before conclusions are drawn in the final section. The optimal regulatory policy is derived for a simple case in Appendix A.

2. Description of the Regulatory Environment

In the regulatory environment considered here, there are three major sources of information asymmetry. The first source arises because the firm's information about the potential for cost reduction, although initially imperfect, is better (in the sense of Blackwell [1951]) than that of the regulator. This phenomenon is captured by the introduction of a random state of nature, θ, the distribution of which is known to the firm, but is not known with certainty by the regulator. It is assumed that the regulator has a non-degenerate prior belief about the actual distribution of θ which is defined over the set of all possible distributions of θ, $\{f^1(\theta),\ldots,f^N(\theta)\}$. Knowledge of the actual realization of θ provides complete information about the potential for cost reduction, since both the regulator and the firm are aware that the maximum value to consumers (i.e., increase in consumers' surplus), x, that can be achieved from any research and development effort, e, aimed at cost reduction in state θ is given by the technology $x = X(e, \theta)$ (the properties of which are defined below).

The second source of information asymmetry centers on the magnitude of research and development (R and D) effort, e. Although the regulator may be able to observe dollar outlays allocated to R and D projects that are ostensibly aimed at cost reduction, he cannot monitor precisely the manner in which these funds are actually employed, nor can he certify the level of intensity or dedication with which the R and D efforts are actually pursued. Consequently, it is assumed here that the regulator cannot observe the actual level of R and D effort put forth by the firm at any point in time, but only the results of such effort, x.

The final source of information asymmetry stems from the fact that after the regulator and firm have "negotiated" (in a sense to be made precise below) a schedule according to which the firm will be compensated for any cost reductions it achieves, the firm's uncertainty about the technology of cost reduction is resolved. Such a resolution of uncertainty might result from an exploratory research phase (or feasibility study) carried out by the firm before it chooses its final level of R and D effort. In terms of the notation introduced above, the feasibility study will permit the firm to determine the actual realization of θ before it chooses e. It is assumed that the study is financed by the regulator as part of the incentive scheme, but the costs of the study (assumed to be known to both the regulator and the firm) are not considered explicitly here. It is conceivable, of course, that under some circumstances the regulated firm might choose to finance its own feasibility study to learn the value of θ before negotiating a compensation schedule with the regulator. Such a possibility is, however, not considered here.[4]

As noted in the introduction, both the firm and the regulator are assumed to have perfect information about the firm's production costs and about the demand functions that face the firm at any moment in time. The regulator is also assumed to know the objective of the firm, which is profit-maximization. In addition, so that the results presented in Section 3 can be attributed solely to differences in information about the technology of cost reduction, all risk-sharing features of the regulatory process are ignored, and both the regulator and the firm are assumed to be risk neutral.

Because there is no uncertainty about the cost and demand functions that prevail in the industry, both the regulator and firm can translate any change in the firm's cost function into the maximum possible increase in consumers' surplus (which is achieved by setting the prices of the firm at their Ramsey levels). Consequently, any compensation schedule that the regulator might offer the firm can, without loss of generality, be expressed as a function of the value of any cost reduction in terms of the associated maximum possible gain in consumers' surplus, x. Furthermore, because the regulator can observe neither e nor θ, any incentive scheme must specify payments to the firm based solely on the single observable variable, x.

In theory, such an incentive scheme may be of any functional form. Here, though, attention is focused on the class of linear compensation schedules which can be expressed as $R(x) = a + bx$, where $R(x)$ are the revenues paid to the firm for achieving cost reductions that result in an increase in consumers' surplus of x dollars, and a and b are constants. In practice, regulatory rules often take simple mathematical forms in order to facilitate their administration and to make their terms and intent apparent to all concerned parties. With a linear incentive scheme, the regulator can report to the public in simple terms the share of any gain in consumers' surplus that will be awarded the firm as an incentive for cost reductions.

It is assumed that the regulator acts faithfully in the interests of the consumers of the firm. More precisely, the objective of the regulator is to maximize the expected increase in consumers' surplus net of payments to the firm. The firm, on the other hand, wishes to maximize the difference between the compensation it receives, $R(x)$, and the costs of the associated level of research and development effort. From the technology of cost reduction specified above and from knowledge of the true costs to the firm of any level of R and D effort, the actual minimum dollar expenditure required to achieve an increase in consumers' surplus of x dollars in state θ can be specified as $D(x,\theta)$. It is plausible to assume that the required expenditures increase (at an increasing rate) with x so that $D_x(x,\theta) \geqq 0$ and $D_{xx}(x,\theta) \geqq 0$ $\forall\theta$, $x \geqq 0$, with strict inequality for $x > 0$. Also, in keeping with the role of θ as a productivity parameter, it is assumed that higher realizations of θ correspond to states in which the costs to the firm of achieving any level of x are smaller, and in which the marginal costs of additional units of x are also smaller, i.e., $D_\theta(x,\theta) \leqq 0$ and $D_{x\theta}(x,\theta) \leqq 0$ $\forall\theta$, $x \geqq 0$ with strict inequality for $x > 0$.

Before the regulator's problem is stated formally, one additional feature of the regulatory environment need be made explicit. The set of possible distributions of θ, $\{f^1(\theta),\ldots,f^N(\theta)\}$, is restricted to any N continuous distributions with strictly positive support on the interval $[\underline{\theta},\overline{\theta}]$ ($\underline{\theta} > 0$) which stand in a relation of first-order stochastic dominance. In other words,

$F^1(\theta) \geq F^2(\theta) \geq \ldots \geq F^N(\theta)\ \forall\theta \in [\underline{\theta},\overline{\theta}]$ where $F^i(\theta) \equiv \int_{\underline{\theta}}^{\theta} f^i(\S)d\S$. Thus, the set of possible distributions might, for example, consist of N exponential distributions or N normal distributions, each with the same variance.

Intuitively, the assumption of stochastic dominance ensures that the higher the numerical value of the superscript on the actual distribution of θ, the more conducive is the technological climate to cost reduction efforts. In other words, an arbitrary research and development effort, e', is more likely to lead to cost reductions that provide an increase in consumers' surplus in excess of any specified level, $x' = X(e',\theta')$, the "higher" the distribution of θ (since $F^i(\theta') \geq F^j(\theta')$ for $i \leq j$).

To reiterate the nature of the informational environment considered here, then, at the time when the regulator designs an incentive scheme to foster cost reductions that would otherwise not be forthcoming, neither the regulator nor the firm has perfect information about the prevailing technology of cost reduction (as captured in θ). Both parties, however, realize that the firm's information is better than that of the regulator. It is also known that before the firm must make a final decision about the magnitude of effort to invest in the R and D process (but after a compensation schedule is agreed upon), the firm, though its feasibility study, will learn exactly how successful such efforts will be (i.e., it will observe the realization of θ).

A straightforward application of the results of Harris and Townsend [1981] reveal that in a situation of this type, the optimal strategy for the less-informed party (the regulator) is to offer the better-informed party (the firm) a choice among contracts (linear compensation schedules). If the set of compensation schedules is designed appropriately, the firm may be induced to use its information to select schedules from among the set offered in a manner consistent with the preferences of the regulator, were he to share the firm's privileged information. Thus, the regulator can do no better than to offer the firm at most N distinct compensation schedules (one corresponding to each of the possible distributions of θ) from which the firm is permitted to make a binding choice.

With this fact in mind, the regulator's problem (RP) can be stated as follows:

$$\underset{R^1(\cdot),\ldots,R^N(\cdot)}{\text{Maximize}} \quad \sum_{i=1}^{N} \phi^i \int [x^i(\theta) - R^i(x^i(\theta))] f^i(\theta) d\theta$$

Subject to: (RP)

(IR) $$\int [R^i(x^i(\theta)) - D(x^i(\theta),\theta)] f^i(\theta) d\theta \geqq 0 \qquad i = 1,\ldots,N$$

(SSB) $$\int [R^i(x^i(\theta)) - D(x^i(\theta),\theta)] f^i(\theta) d\theta \geqq \int [R^j(x^j(\theta)) - D(x^j(\theta),\theta)] f^i(\theta) d\theta \qquad \forall i, j = 1,\ldots,N$$

(SSW) $$x^i(\theta) = \underset{x'}{\text{argmax}}\ R^i(x') - D(x',\theta) \qquad \forall i = 1,\ldots,N$$

where $R^i(x) = a^i + b^i x$ is the compensation scheme that the firm will select when the actual distribution of θ is $f^i(\theta)$,

$x^i(\theta)$ = the increase in consumers' surplus that will be forthcoming in state θ under compensation scheme $R^i(\cdot)$,

ϕ^i = probability that $f^i(\theta)$ is the actual distribution of θ,

and where all integrals are defined here (and elsewhere unless otherwise specified) over the interval $[\underline{\theta},\overline{\theta}]$.

The individual rationality (IR) constraints indicate that the firm will only agree to the terms a compensation scheme that it (weakly) prefers to the absence of any scheme. Lacking precise information about the technology of cost reduction and not wishing to impose a policy that would bankrupt the firm, the regulator will not compel the firm to accept any scheme that, under the actual distribution of θ, promises an expected loss to the firm. The (SSB) constraints guarantee self-selection between contracts, i.e., they ensure that the firm will choose compensation scheme $R^i(x)$ when $f^i(\theta)$ is the actual distribution of θ. The (SSW) constraints ensure self-selection within each contract, i.e., they guarantee that the firm will undertake sufficient research and development in state θ to realize an increase in consumers' surplus of $x^i(\theta)$ under compensation scheme $R^i(x)$. It is assumed throughout that when indifferent among two or more incentive schemes (or levels of effort), the firm will choose the one preferred by the regulator.

3. Properties of the Optimal Incentive Scheme

Absent imperfect and asymmetric information, the regulator would simply order the firm to undertake the socially optimal level of cost reduction, i.e., to reduce costs sufficiently so as to increase consumers' surplus to the point at which further increases would impose costs upon the firm that exceed the gain in surplus to consumers. In terms of the notation employed here, the socially optimal level of increase in consumers' surplus in state θ_i, $x^*(\theta_i)$, is determined by the equation $D_x(x^*(\theta_i),\theta_i) = 1$.

Similarly, if the set of possible distributions were narrowed to a single distribution, $f^S(\theta)$, so that the regulator and firm shared the same information ex ante, the optimal strategy for the regulator would be to offer the firm the single compensation schedule $R^S(x) = a^{*S} + x$ where $a^{*S} \equiv - \int [x^*(\theta) - D(x^*(\theta),\theta)]f^S(\theta)d\theta$. This scheme would be efficient in that it would induce the firm to realize the increase in consumers' surplus that is ex post Pareto efficient in the particular state of nature that is eventually realized, whatever that state might be.[5] The scheme also extracts all of the firm's expected gain, so that $R^S(x)$ is clearly the best scheme from the point of view of the regulator.[6] It is also the best scheme from a social viewpoint (i.e., one in which the gains of consumers and the firm are equally weighted) since it is an efficient scheme.

It is also true that the regulator can guarantee an efficient outcome in every state of nature even under the more general conditions considered here. To do so, he need only offer the firm the compensation scheme $R(x) = a^1 + x$ where $a^1 = -\int [x^*(\theta) - D(x^*(\theta),\theta)]f^1(\theta)d\theta$ (< 0 by assumption). This scheme, however, will grant the firm a strictly positive expected return under every distribution except $f^1(\theta)$. (Recall that $f^1(\cdot)$ is stochastically dominated by all other possible distributions.) Also, the gain to consumers from this scheme is identically $|a^1|$ whatever the actual distribution of θ. Whether such a strategy is ever optimal for the regulator to pursue is addressed in Proposition 1. First, though, Lemmas 1 through 4 describe some findings that are essential to the proof of the Proposition, and also provide some insight into the optimal set of compensation schedules.

Lemma 1. The increase in consumers' surplus that will be realized under any compensation scheme $R(x) = a + bx$ with b (strictly) positive is a (strictly increasing function of θ.

Proof. Let $x(b,\theta)$ represent this increase. By (SSW) in (RP), $x(b,\theta)$ is determined by the equation $D_x(x(b,\theta),\theta) = b$.

Total differentiation with respect to θ reveals that $D_{xx}(\cdot,\cdot)x_\theta + D_{x\theta}(\cdot,\cdot) = 0$, which implies that $x_\theta(\cdot,\cdot)$ is non-negative since, by assumption, $D_{xx}(\cdot,\cdot) \geqq 0$ and $D_{x\theta}(\cdot,\cdot) \leqq 0$. $x_\theta(\cdot,\cdot)$ is strictly positive whenever $x(\cdot)$ is strictly positive, which will always be the case if $b > 0$.

Q.E.D.

Lemma 2. The increase in consumers' surplus that will be realized in any state θ is a strictly increasing function of the slope, b, of the compensation schedule selected by the firm, $\forall b > 0$.

Proof. Using the notation developed in Lemma 1, it follows from (SSW) in (RP) that $D_x(x(b,\theta),\theta) = b$.

Hence, $D_{xx}(\cdot,\cdot)x_b(\cdot) = 1$, so that since $D_{xx}(\cdot,\cdot) > 0 \; \forall b > 0$, $x_b(\cdot) > 0$ also.

Q.E.D.

Lemma 3. The expected net return, to the firm under any compensation schedule, $\int [R(x) - D(x,\theta)]f^i(\theta)d\theta$, is increasing in i (i.e., the firm's expected return is greater the "more productive" the environment).

Proof. Under distribution $f^i(\theta)$, the firm's expected net return from schedule $R(x) = a + bx$ can, using the notation of Lemma 1, be written as $\int [bx(b,\theta) - D(x(b,\theta),\theta) + a]f^i(\theta)d\theta$. Since $D_x(\cdot,\cdot) = b$ and $D_\theta(\cdot,\cdot) < 0$, $[bx(b,\theta) - D(x(b,\theta),\theta)]$ is an increasing function of θ. Hence, the Lemma follows from the assumption that $F^i(\theta) \geqq F^j(\theta) \; \forall \theta \in [\underline{\theta},\overline{\theta}]$ whenever $i < j$.

Q.E.D.

Lemma 4. Suppose that the optimal set of compensation schedules (i.e., the solution to (RP)) leaves the firm indifferent between $R^i(x)$ and $R^j(x)$ under either distribution $f^i(\theta)$ or $f^j(\theta)$. Then this set cannot contain another schedule $R^h(x)$ that is distinct from $R^i(x)$ and $R^j(x)$ and (weakly) preferred to both by the firm under distribution $f^h(\theta)$, where $i < h < j$.

Proof. The proof proceeds by contradiction and is discussed here only for the case in which the firm is indifferent between $R^i(x)$ and $R^j(x)$ under distribution $f^j(\theta)$. The corresponding proof under $f^i(\theta)$ is similar, and is omitte

Since $R^h(x)$ is preferred by the firm to $R^i(x)$ under $f^h(\theta)$, $\int g(\theta;b^h,b^i)f^h(\theta)d\theta \geqq 0$ from (SSB) in (RP), where $g(\theta;b^h,b^i) = [b^h x(b^h,\theta) - D(x(b^h,\theta),\theta) + a^h] - [b^i x(b^i,\theta) - D(x(b^i,\theta),\theta) + a^i]$. It is straightforward to verify that $g(\theta;b^h,b^i)$ is an increasing function of θ when $b^h > b^i$, and that $g(\cdot)$ also increases as its second argument increases. Consider the case in which $b^j > b^h > b$

$$\int g(\theta;b^j,b^i)f^j(\theta)d\theta > \int g(\theta;b^j,b^i)f^h(\theta)d\theta$$
$$> \int g(\theta;b^h,b^i)f^h(\theta)d\theta \geqq 0.$$

The first inequality holds because $F^h(\theta) \geqq F^j(\theta)$ $\forall\theta \in [\underline{\theta},\overline{\theta}]$. The second inequality holds because $g(\theta;b^j,b^i) > g(\theta;b^h,b^i)$ $\forall\theta \in [\underline{\theta},\overline{\theta}]$. This result states that the firm strictly prefers $R^j(x)$ to $R^i(x)$ under distribution $f^j(\theta)$, which violates the hypothesis.

Now consider the case in which $b^i = b^j$. If $b^h > b^j = b^i$, then

$$\int g(\theta;b^h,b^j)f^j(\theta)d\theta > \int g(\theta;b^h,b^j)f^h(\theta) \geqq 0$$

which is a contradiction. And if $b^h < b^j = b^i$, then

$$\int g(\theta;b^h,b^i)f^i(\theta)d\theta > \int g(\theta;b^h,b^i)f^h(\theta)d\theta \geqq 0$$

because $g(\theta;b^h,b^i)$ is a decreasing function of θ when $b^h < b^i$. Hence, another contradiction is reached. Arguments presented in the proof of Proposition 1 reveal that other cases need not be considered as possible solutions to (RP).

Q.E.D.

Lemmas 1 through 4 are instrumental in the proof of Proposition 1. This proposition describes the major features of the optimal regulatory plan to foster research and development. It is apparent that the optimal plan will not ensure that an ex post Pareto efficient outcome is always realized, even though it is within the regulator's powers to do just that. In the statement of the proposition, $E\pi(R^i|f^j)$ represents the expected net return to the firm under incentive scheme $R^i(x)$ when $f^j(\theta)$ is the actual distribution of θ.

Proposition 1. The optimal set of linear compensation schedules $\{R^i(x) = a^i + b^i x$

$i = 1,\ldots,N\}$ has the following properties:

i) $a^1 \geqq a^2 \geqq \ldots \geqq a^{N-1} > a^N$

ii) $0 \leqq b^1 \leqq \ldots \leqq b^{N-1} < b^N$

iii) $b^N = 1$

iv) $E\pi(R^i|f^i) = 0 \quad \forall i = 1,\ldots,\bar{i}$, where $\bar{i}$ is the smallest value of i for which $b^i > 0$,

v) $E\pi(R^i|f^i)$ is increasing in i, and strictly so if $b^i > 0$

vi) $E\pi(R^i|f^i) = E\pi(R^{i-1}|f^i) \quad \forall i = 2,\ldots,N$.

Proof of Proposition 1.

Let $\bar{R}^i(x) = \bar{a}^i + \bar{b}^i x$ be the $M \leqq N$ distinct compensation schedules with strictly positive slope that (perhaps along with $\bar{R}(x) = 0$) comprise the optimal set. Note that it is without loss of generality that compensation schedules with strictly negative slope are not considered in the solution to (RP), because such schedules will always induce the regulated firm to select $x = 0$, as will a schedule with slope of zero.

Furthermore, let I^i be the set of all superscripts on those distributions under which the firm will choose $\bar{R}^i(x)$. Then (RP) can be rewritten as

$$\underset{\overline{a}^i,\overline{b}^i}{\text{Maximize}} \quad \sum_{i=1}^{M} \sum_{h\varepsilon I^i} \phi^h\{(1-\overline{b}^i)\int x(\overline{b}^i,\theta)f^h(\theta)d\theta - \overline{a}^i\}$$

subject to: (RP')

(IR')
$$\int [\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i]f^h(\theta)d\theta \geqq 0$$
$$\forall h \varepsilon I^i, \qquad i = 1,\ldots,M$$

(SSB')
$$\int [\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i]f^h(\theta)d\theta \geqq$$
$$\int [\overline{b}^j x(\overline{b}^j,\theta) - D(x(\overline{b}^j,\theta),\theta) + \overline{a}^j]f^h(\theta)d\theta$$
$$\forall h \varepsilon I^i, \qquad i,j = 1,\ldots,M$$

where, as defined in Lemma 1, $x(\overline{b}^i,\theta)$ is the increase in consumers' surplus that the firm will produce in state θ after choosing compensation scheme $\overline{R}^i(x)$.

The Lagrangian function associated with (RP') is

$$\begin{aligned} L = & \sum_{i=1}^{M} \sum_{h\varepsilon I^i} \phi^h\{(1-\overline{b}^i)\int x(\overline{b}^i,\theta)f^h(\theta)d\theta - \overline{a}^i\} \\ & + \sum_{i=1}^{M} \sum_{h\varepsilon I^i} \lambda^{ih}\int [\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i]f^h(\theta)d\theta \\ & + \sum_{i=1}^{M} \sum_{\substack{j=1 \\ j\neq i}}^{M} \sum_{h\varepsilon I^i} \gamma^{ijh}\int \{[\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i] \\ & \qquad - [\overline{b}^j x(\overline{b}^j,\theta) - D(x(\overline{b}^j,\theta),\theta) + \overline{a}^j]\}f^h(\theta)d\theta. \end{aligned}$$

To simplify this expression, it is helpful to determine which of the multipliers, if any, are identically zero.

First, from (SSB'), Lemma 3, and (IR'), it follows that

$$\int [\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i] f^h(\theta) d\theta$$

$$\geqq \int [\overline{b}^1 x(\overline{b}^1,\theta) - D(x(\overline{b}^1,\theta),\theta) + \overline{a}^1] f^h(\theta) d\theta$$

$$\forall h \in I^i \quad \text{for each} \quad i = 2,\ldots,M$$

$$> \int [\overline{b}^1 x(\overline{b}^1,\theta) - D(x(\overline{b}^1,\theta),\theta) + \overline{a}^1] f^k(\theta) d\theta \geqq 0$$

$$\forall k \in I^1.$$

Hence, $\lambda^{ih} = 0 \quad \forall h \in I^i$ for each $i = 2,\ldots,M$.

Second, defining $n(i) \equiv \text{minimum } \{I^i\}$, it follows from Lemma 3 that $\lambda^{1h} = 0$ $\forall h \neq n(1)$.

Third, it follows directly from Lemma 4 that $\gamma^{ijh} = 0 \quad \forall j > i+1$ and $\forall j < i-1$, for each $h \in I^i$.

Fourth, if $\gamma^{i,i-1,h} > 0$ for some $h > n(i)$ and $h \in I^i$, then by Lemma 3, $E\pi(\overline{R}^i | f^h) = E\pi(\overline{R}^{i-1} | f^h) > E\pi(\overline{R}^{i-1} | f^{n(i)})$ which violates the definition of $\overline{R}^i(x)$ as the compensation schedule that will be selected by the firm under all distribution $f^h(\theta)$ where $h \in I^i$. Hence, $\gamma^{i,i-1,h} = 0 \quad \forall h \neq n(i)$ and $h \in I^i$, for each $i = 2,\ldots,M$.

A similar argument by contradiction explains why $\gamma^{i,i+1,h} = 0 \quad \forall h < m(i)$ $\equiv$ maximum $\{I^i\}$ and $h \in I^i$, for each $i = 1,\ldots,M-1$.

Using these results and differentiating the Lagrangian function with respect to $\overline{a}^i$, M of the necessary conditions for a maximum reduce to

$$\sum_{h\varepsilon I^i}\phi^h - \lambda^{i,n(i)} - \gamma^{i,i-1,n(i)} - \gamma^{i,i+1,m(i)} + \gamma^{i-1,i,m(i-1)}$$

$$+ \gamma^{i+1,i,n(i+1)} = 0 \qquad (3.1)$$

$$\forall\ i = 1,\ldots,M$$

where $\gamma^{1,0,h} = \gamma^{M,M+1,h} = \gamma^{0,1,h} = \gamma^{M+1,M,h} = 0 \quad \forall\ h.$

Summing all of these equations reveals that $\lambda^{1,n(1)} = 1$, which along with Lemma 3, proves property (iv) of the Proposition. (Note that property (v) is an immediate consequence of Lemma 3 and the (SSB') constraints.) Then, from equation (3.1) with $M = 1$, it follows that

$$\gamma^{2,1,n(2)} = 1 - \sum_{h\varepsilon I^1}\phi^h + \gamma^{1,2,m(1)} > 0.$$

A proof by contradiction which proceeds, much as does the proof of Lemma 2 reveals that whenever $\gamma^{i,i-1,n(i)} > 0$, $\gamma^{i-1,i,m(i-1)} = 0$. Hence, $\gamma^{2,1,n(2)} = 1 - \sum_{h\varepsilon I^1}\phi^h$. A straightforward induction proof that employs these two findings and equation (3.1) reveals that $\gamma^{i,i-1,n(i)} = 1 - \sum_{j=1}^{i=1}\sum_{h\varepsilon I^j}\phi^h > 0 \quad \forall\ i = 2,\ldots,M$ and therefore, $\gamma^{ijh} = 0$ for all other i, j, h. Hence, by the complementary slackness conditions associated with (SSB'), property (vi) of the Proposition is proved. It only remains, then, to prove the first three properties. The foregoing findings allow simplification to the following equations of the M necessary conditions for a maximum derived from differentiating the Lagrangian with respect to $\bar{b}^i$:

$$(1-\overline{b}^i) \sum_{h\varepsilon I^i} \phi^h \int x_b(\overline{b}^i,\theta) f^h(\theta) d\theta$$

$$= [1 - \sum_{j=1}^{i} \sum_{h\varepsilon I^j} \phi^h] \int x(\overline{b}^i,\theta)[f^{n(i+1)}(\theta) - f^{n(i)}(\theta)]d\theta$$

$$+ \sum_{h\varepsilon I^i} \phi^h \int x(\overline{b}^i,\theta)[f^h(\theta) - f^{n(i)}(\theta)]d\theta \qquad (3.2)$$

$$\forall\ i = 1,\ldots,M.$$

where $f^{n(M+1)}(\theta) = 0 \quad \forall\ \theta$.

By Lemma 1 and the assumption of stochastic dominance, the first term on the right-hand side of (3.2) is strictly positive for all $i = 1,\ldots,M-1$ and zero for $i = M$; the second term is non-negative $\forall\ i$, and is strictly positive whenever $\overline{R}^i(x)$ is selected by the firm under more than one distribution of θ. Consequently, $\overline{b}^i < 1 \quad \forall\ i = 1,\ldots,M-1$ since $x_b(\overline{b}^i,\theta) > 0 \ \forall\ \theta$ by Lemma 2. Furthermore, $\overline{b}^M \leq 1$, with equality if and only if $I^M = \{N\}$. It is proved in Appendix C that the regulator will never induce the firm to choose a contract whose slope is strictly less than unity when $f^N(\theta)$ is the true distribution of θ. Thus, $\overline{b}^M = b^N = 1$.

It only remains to be shown, then, that $\overline{b}^1 < \ldots < \overline{b}^{M-1}$ which, along with Lemma 4, is sufficient to prove property (ii) in Proposition 1. Property (i) follows immediately from property (ii) since if a particular compensation schedule had both a smaller slope and a smaller intercept than another schedule, the former would never be selected by the firm.

Suppose $\overline{b}^i > \overline{b}^j$ for some $i < j$. From the (SSB') constraints, $\int g(\theta;\overline{b}^i,\overline{b}^j) f^{n(i)}(\theta) > 0$ where, as defined above,

$$g(\theta,\overline{b}^i,\overline{b}^j) = [\overline{b}^i x(\overline{b}^i,\theta) - D(x(\overline{b}^i,\theta),\theta) + \overline{a}^i]$$

$$- [\overline{b}^j x(\overline{b}^j,\theta) - D(x(\overline{b}^j,\theta),\theta) + \overline{a}^j].$$

As noted in Lemma 4, $g_\theta(\theta;\overline{b}^i,\overline{b}^j) > 0$ for $\overline{b}^i > \overline{b}^j$, so that since $F^i(\theta) \geq F^j(\theta) \quad \forall\, \theta \in [\underline{\theta},\overline{\theta}]$, $\int g(\theta;\overline{b}^i,\overline{b}^j)f^{n(j)}(\theta)d\theta > \int g(\theta;\overline{b}^i,\overline{b}^j)f^{n(i)}(\theta)d\theta \geq 0$.

This result states that the firm strictly prefers compensation schedule $\overline{R}^i(x)$ to $\overline{R}^j(x)$ when $f^{n(j)}(\theta)$ is the actual distribution of θ, which violates a (SSB') constraint. Hence, for $i < j$, $\overline{b}^i < \overline{b}^j$.

Q.E.D.

There are a few features of the optimal regulatory strategy that deserve particular emphasis. First, the sequential ordering of slopes (and intercepts) of the compensation schedules is of a particularly simple nature. Second, the fact that the slopes of the compensation schedules will always lie between zero and unity makes the interpretation of the slope as a "sharing ratio" particularly appealing. Third, it is only in those environments least conducive to cost reduction that the firm does not retain any of the additional surplus generated by its efforts. And fourth, it is only in the most productive environment, $f^N(\theta)$, that the increase in consumers' surplus that is actually realized in any state is the amount that is ex post Pareto efficient in that state. A complete explanation for these features of the optimal regulatory strategy is presented in Section 4. The discussion in Appendix A is presented in order to provide a feel for the magnitude of the relevant values in a particular example.

4. An Explanation of the Finding

When he designs an incentive scheme to promote research and development aimed at cost reduction, the regulator must be concerned both with the total expected increase in surplus generated by any compensation schedule and the share of this total that will be awarded consumers. Ideally, the regulator would like to award consumers the largest possible share of the greatest possible expected total increase in surplus regardless of how conducive the environment may be to successful cost reduction. In most regulatory situations, though, it will not be possible for the regulator to do so because of uncertainty about the technology of cost reduction.

As noted in Section 3, the regulator can, despite the informational limitations discussed there, always ensure that the total expected increase in surplus, $\int [x(\theta) - D(x(\theta),\theta)]f^i(\theta)d\theta$, is as large as possible under every possible distribution, $f^i(\theta)$, by offering the firm the single efficient schedule $R^1(x) = a^{*1} + x$ where $a^{*1} \equiv -\int [x^*(\theta) - D(x^*(\theta),\theta)]f^1(\theta)d\theta$. In doing so, however, the regulator salvages all of the expected gain in surplus for consumers only in the environment least conducive to cost reduction (i.e., only when $f^1(\theta)$ is the actual distribution of θ), and awards "too large" a share of the total surplus to the firm under all other situations. The sense in which the firm's share is too large can be made precise in the following manner.

Recall from Lemma 3 that the level of expected net payoff to the firm from any feasible compensation scheme (i.e., any one that satisfies (SSW) in (RP)) is greater under $f^j(\theta)$ than under $f^i(\theta)$ whenever $i < j$. Consequently, if the firm is granted an expected net payoff of, say, π^i by the optimal set of compensation schedules under distribution $f^i(\theta)$, it must receive strictly more than π^i under distribution $f^j(\theta)$ since, at worst, the firm can always choose the same schedule it would choose under distribution $f^i(\theta)$. The additional expected payoff that must be awarded the firm in the more productive environment, though, will not exhaust the entire gain in surplus afforded by the occurence of $f^j(\theta)$ rather than $f^i(\theta)$ as long as the slope of the compensation schedule selected by the firm under $f^i(\theta)$ is less than unity. Thus, as long as the slope of each compensation schedule in the optimal set does not exceed unity, the regulator can always salvage for consumers some fraction of any incremental surplus generated under an environment more conducive to research and development.

Furthermore, the closer is the slope of some compensation schedule, $R^i(x)$, to unity, the smaller is the fraction of any incremental surplus that the regulator can retain for consumers in the event of distributions that stochastically dominate the one in which $R^i(x)$ is selected by the firm. Thus, although a compensation schedule with a slope (b^i) less than unity is inefficient in the sense that it does not maximize the total expected surplus (to consumers and the firm) under the distribution, $f^i(\theta)$, in which it is selected, such a schedule is consistent with the regulator's desire to maximize the expected increase in consumers' surplus because it reduces the share of incremental surplus that must be awarded the firm under distributions that stochastically dominate $f^i(\theta)$ below the corresponding share if b^i were (greater than or) equal to unity.

In order to formally prove the foregoing assertions, it is helpful to introduce the following terminology:

<u>Total Incremental Surplus (IS)</u> $\equiv \int [x^*(\theta) - D(x^*(\theta),\theta)]f^j(\theta)d\theta$

$$- \int [x(b^i,\theta) - D(x(b^i,\theta),\theta)]f^i(\theta)d\theta$$

is the maximum possible increase in total expected surplus in excess of the level achieved under $f^i(\theta)$ with compensation schedule $R^i(x) = a^i + b^i x$ that can be achieved under distribution $f^j(\theta)$ $(j > i)$.

Minimum Incremental Surplus to the Firm (MISF) $\equiv \int [b^i x(b^i,\theta) - D(x(b^i,\theta),\theta)][f^j(\theta) - f^i(\theta)]d\theta$ is the minimum amount of additional surplus that must be awarded the firm under $f^j(\theta)$ above the amount awarded the firm by schedule $R^i(x)$ under $f^i(\theta)$.

Fraction of Incremental Surplus to the Firm (FISF) $\equiv$ MISF/IS.

The expression for MISF cited above follows directly from the (SSB) constraints in (RP). Note that if the regulator induces the firm to choose an efficient compensation schedule under distribution $f^i(\theta)$, then IS = MISF so that FISF = 1. In other words, if the optimal set of compensation schedules is such that the firm chooses an efficient schedule under some distribution, $f^i(\theta)$ $i \neq N$, then the firm must be awarded all of the additional expected gain that is generated under every distribution that stochastically dominates $f^i(\theta)$. At the other extreme, if $b^i = 0$, then FISF = 0 and the regulator can retain the entire gain in expected surplus for consumers. Furthermore, it is straightforward to verify that the expression for IS decreases and that for MISF increases with increases in b^i, so that as b^i approaches unity (from below), the value of FISF increases towards unity. It can also be shown that MISF > IS whenever $b^i > 1$, so that FISF > 1 (and FISF increases as b^i increases above unity). In other words, if the schedule selected by the firm under distribution $f^i(\theta)$ has a slope greater than unity, then consumers will be forced to expect less consumers' surplus under all distributions which stochastically dominate $f^i(\theta)$ than they receive under $f^i(\theta)$, even though a greater total expected surplus is always available under these other distributions.

This lattermost conclusion explains why it will never be optimal for the regulator to set in excess of unity the slope of any compensation schedule that will be selected by the firm. Also, the regulator will never want to set at unity the slope of any schedule other than $R^N(x)$ since by doing so, he would foreclose the opportunity for consumers to share in the larger expected gains in surplus under technological environments more conducive to research and development.

The extent to which the slopes of all schedules other than $R^N(x)$ will fall short of unity depends upon the relative frequency with which the various distributions occur and the differences in total expected surplus gains among the distributions. The smaller the slope of any schedule, the smaller is the fraction of the potential total expected increase in surplus that is actually realized when the particular distribution occurs under which this schedule is selected. Offsetting this effect, though, is the aforementioned fact that the smaller the slope of any schedule, the greater the fraction of additional expected surplus generated in environments more conducive to cost reduction that can be retained for consumers. Hence, as the comparative static results in Appendix A suggest, the regulator will be more willing to forego gains in surplus under some distribution, $f^i(\theta)$, the more probable are distributions that stochastically dominate $f^i(\theta)$ thought to be and the less probable is $f^i(\theta)$, itself, thought to be.

5. Conclusions

In the presence of perfect information about costs and demand, the optimal strategy for a regulator charged with maximizing consumers' surplus (or social welfare) in any one period is well-documented (in, for example, Baumol and Bradford [1970]). Recent advances have also been made in describing the optimal strategy for a regulator in the presence of asymmetric information about current costs and/or demand (see, for example, Baron and DeBondt [1981], Baron and Myerson [1982], Loeb and Magat [1979], Sappington [1982] and Vogelsang and Finsinger [1979]). Little attention, though, has heretofore been afforded the optimal strategy of a regulator charged with maximizing the present value of consumers' surplus when the regulated firm has the unique ability to effect cost reductions in its operations and when the firm has more precise information about the technology of cost reduction than does the regulator.

This research has analyzed the optimal strategy for a regulator under such conditions. It was shown that the regulator should design a set of compensation schedules from which the firm can employ its superior information to make a binding choice. The optimal set of linear schedules are such that the firm will expect only to break even in those environments least conducive to cost reduction. In each of the "more productive" environments, the firm will select more profitable schedules that permit it to retain successively larger fractions (not in excess of unity) of any increase in consumers' surplus that its efforts generate, less some fixed payment. Thus, in order to induce the firm to employ its privileged information and abilities in the process of cost reduction, the regulator should not, in general, limit the firm to zero profits as he should do in a world of certainty where questions of moral hazard and adverse selection do not arise.

Needless to say, there are a number of important ways in which the model developed here abstracts from reality. First, regulation is an ongoing process in which the regulator and firm interact and acquire information about each other over time. Intertemporal strategic reaction to regulatory incentives, therefore, may be important to consider. Second, the regulator's information about actual costs and demand is likely to be imperfect. Third, both the regulator and firm may be extremely averse to situations in which the firm's overall profits are negative (situations that are not ruled out by the incentive scheme developed above). And fourth, the range of environments which reflect the potential for cost reduction may not be so easily identified and ranked by the regulator. These considerations remain as topics for future research.

Appendix A

The optimal set of compensation schedules is derived here for the case in which θ may follow one of three possible uniform distributions,[7] i.e., $f^i(\theta) \sim U(\underline{\theta},\theta^i)$ $i = 1,2,3$ with $\underline{\theta} < \theta^1 < \theta^2 < \theta^3$. At the outset, the firm knows the actual distribution of θ while the regulator believes that the actual distribution is $f^i(\theta)$ with probability $\phi^i > 0$, $i = 1,2,3$. It is also assumed here that $D(x,\theta) = 2x^2/\theta$. The solution to (RP) under these assumptions may consist of either two or three distinct compensation schedules. Attention is focused here on the latter case.[8] When the values of ϕ^i and θ^i $i = 1,2,3$ are such that the regulator will design three distinct schedules, the slopes of these schedules will be:

$$b^1 = \phi^1\theta^1/[\phi^1\theta^1 + (\phi^2 + \phi^3)(\theta^2 - \theta^1)] ,$$

$$b^2 = \phi^2\theta^2/[\phi^2\theta^2 + \phi^3(\theta^3 - \theta^2)] , \text{ and } b^3 = 1.$$

The corresponding values for a^1, a^2, and a^3 can be computed and shown to satisfy properties (iv) and (vi) in Proposition 1.

Straightforward manipulation of these results reveals that the slope of the $R^i(x)$ schedule $i = 1,2$ increases as ϕ^i increases and $\sum_{j=i}^{3} \phi^j$ decreases by the same amount. b^i also increases as θ^{i+1} decreases. Thus, as suggested by the arguments in Section 4, the regulator will be less willing to forego any available surplus in the environment represented by $f^i(\theta)$: (i) the more sure he is that this environment, rather than more productive ones, is the actual one, and (ii) the smaller the amount by which the total expected surplus from any schedule under $f^{i+1}(\theta)$ exceeds that under $f^i(\theta)$.

Figure 1 illustrates the optimal triple of compensation schedules for the particular case in which the regulator's prior beliefs are diffuse, i.e., $\phi^1 = \phi^2 = \phi^3 = 1/3$, and in which $\underline{\theta} = 0$, $\theta^1 = 1$, $\theta^2 = 2$, and $\theta^3 = 4$. In this case, the three schedules are:

$$R^1(x) = -2/288 + 1/3x$$

$$R^2(x) = -7/288 + 1/2x$$

$$R^3(x) = -61/288 + 1x \ .$$

Figure 1 also illustrates how a particular value of x is selected by the firm after a compensation schedule, say $R^1(x)$, is chosen. Subsequent to the actual realization of θ, say $\hat{\theta}$, the firm will choose from among the feasible (x,R) - pairs defined by $R^1(x)$ the one $(\overline{x},\overline{R})$ that provides the greatest net return in state $\hat{\theta}$. The curve labelled $\pi(\hat{\theta})$ is Figure 1 depicts those (x,R) pairs among which the firm is indifferent in state $\hat{\theta}$. It's equation is given by $R(x) - 2x^2/\hat{\theta} = k$, where k is a constant. Note that net returns to the firm increase with movements in a northwesterly direction.

Appendix B

A brief discussion is offered here of how the conclusions drawn above would differ if the firm's information about its technological environment were perfect before the optimal regulatory policy was formulated.

If the firm knew the actual realization of θ from the outset, then it would never face any uncertainty about its technological environment. Possession of such information would alter the nature of the optimal regulatory policy to promote research and development in a number of fundamental ways. First, the firm would no longer accept any scheme that, ex post, caused it to suffer a loss in profits. This is not the case under the scenario considered above, since the risk neutral firm was willing to accept a schedule that increased its expected (ex ante) profits, but not necessarily its ex post profits. Analytically, when the firm knows θ from the outset, the individual rationality (IR) constraints must be rewritten as ex post rather than ex ante constraints.

Second, if the firm knows that the actual realization of θ is, say, $\hat{\theta}$, then there is no longer any advantage to the regulator of having the firm first choose a compensation schedule from among a particular set and then decide upon a level of research and development expenditures. Faced with such a "dual" choice, the firm would always choose the schedule that offered the greatest net return in the only relevant state, $\hat{\theta}$, and then proceed to realize that net return. Thus, the optimal regulatory policy would have the regulator offer the firm only one compensation schedule. This optimal schedule would maximize the expected net gain in consumers' surplus while respecting the firm's right not to accept any arrangement that would result in a loss of profits.

The mathematical structure of such a model would correspond closely to that developed by Sappington [1982]. Using the techniques that he develops there, it is straightforward to prove that when the firm has complete information from the start, the optimal linear incentive schedule will have a slope whose magnitude is strictly between zero and unity. Thus, an outcome that is ex post Pareto efficient will never be induced, no matter what the actual realization of θ.

Appendix C

It is proved here that the regulator will never induce the firm to choose a contract whose slope is strictly less than unity when $f^N(\theta)$ is the true distribution of θ, i.e., $b^N = 1$.

Suppose $\vec{b}^M = b^N < 1$, and the corresponding optimal set of compensation schedules is $\{\vec{R}^1(x), \ldots, \vec{R}^M(x)\}$.

Then, a set of compensation schedules that provides a strictly greater value for the regulators objective function in (RP) is $\{\vec{R}^1(x), \ldots, \vec{R}^M(x), \overline{R}^*(x)\}$ where

$$\overline{R}^*(x) = \overline{a}^{*N} - \vec{a}^{MN} + x$$

and where $$a^{*N} \equiv -\int [x^*(\theta) - D(x^*(\theta),\theta)]f^N(\theta)d\theta ,$$

and $$\vec{a}^{MN} \equiv -\int [b^M x(b^M,\theta) - D(x(b^M,\theta),\theta) + a^M]f^N(\theta)d\theta .$$

$\overline{R}^*(x)$ is designed so that: (i) the firm is indifferent between it and $\vec{R}^M(x)$ under $f^N(\theta)$, (ii) it induces an ex post Pareto efficient outcome in each state since $\overline{b}^* = 1$, and (iii) it will not be selected by the firm under any distribution other than $f^N(\theta)$. Hence, the only effect of introducing $\overline{R}^*(x)$ is to increase the total expected surplus under $f^N(\theta)$ without increasing the firm's share of this larger total. Therefore, the regulator's expected payoff is increased through the introduction of $\overline{R}^*(x)$.

Q.E.D.

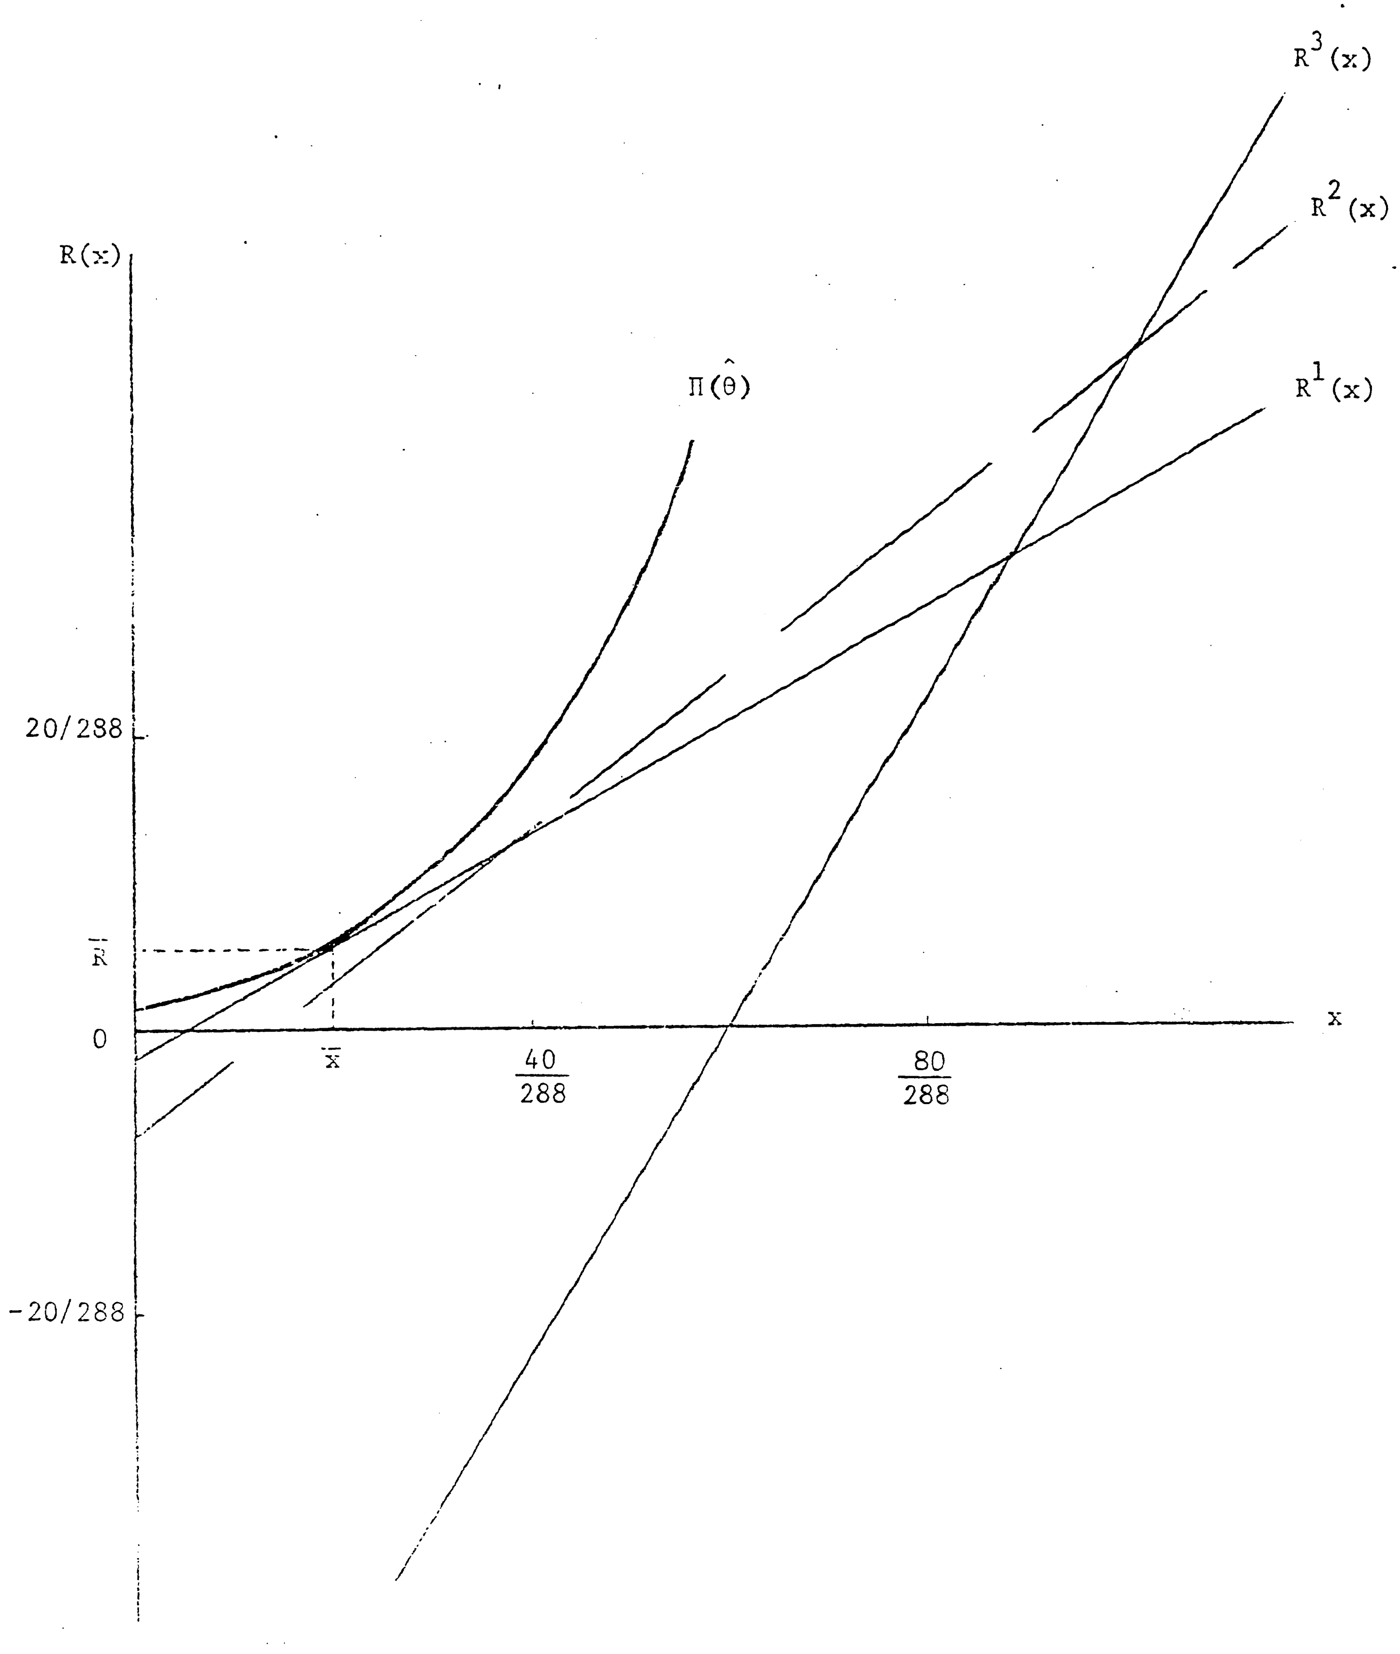

Figure 1. The Optimal Set of Compensation Schedules $\{R^1(x), R^2(x), R^3(x)\}$

FOOTNOTES

1. It should be noted that although the analysis proceeds here in terms of incentive schemes designed to promote cost reduction, the results derived also apply, with some minor re-interpretation, to incentive schemes designed to foster improvements in the quality of the firm's products.

2. In the ensuing analysis, the regulator is assumed to base the firm's compensation on the realized gain in consumers' surplus, x. The variable x can, however, be interpreted more broadly; for example, x could represent the sum of consumers' surplus and profit.

3. The restriction to linear incentive schemes also facilitates comparison of the results derived here with others in the literature. See, for example, Weitzman [1980].

4. If the firm faces no uncertainty about the technology of cost reduction from the outset and the regulator is aware of this fact, then the relevant incentive problem differs fundamentally from that considered here. For one thing, it is possible to show that the expected payoff to consumers is strictly greater when the firm commits itself to a binding compensation scheme before, rather than after, it learns the true state of nature. Thus, the regulator charged with maximizing consumers' surplus will, whenever possible, negotiate with the firm before the latter learns the true value of θ. Other ways in which the incentive problem differs if the firm knows θ from the outset are discussed in Appendix B.

5. Such behavior is induced because under $R^S(x)$, the firm will effectively choose x to maximize the quantity $\{x - D(x,\theta) + a^{*S}\}$. $x^*(\theta)$ is the value of x that maximizes this expression.

6. This result follows directly from findings in the principal-agent literature. See, for example, Harris and Raviv [1979], Holmstrom [1978], and Shavell [1979].

7. Note that this example differs from the standard problem posed in (RP) because the support of each distribution here is not identical. As the results which follow demonstrate, though, all of the properties of the regulator's optimal strategy described in Proposition 1 also characterize the solution to the problem considered here.

8. A sufficient condition for the solution to (RP) to consist of only two distinct contracts in this example is

$$\frac{\phi^1\phi^3}{\phi^2(\phi^2 + \phi^3)} \geqq \frac{\theta^2(\theta^2 - \theta^1)}{\theta^1(\theta^3 - \theta^2)} \tag{A1}$$

This condition is derived by proving that a contradiction is reached if (A1) holds and the solution to (RP) is assumed to consist of three distinct schedules. When the parameters of the problem are such that (A1) does not hold, the optimal strategy for the regulator can be determined through a comparison of his expected payoff under the optimal pair and triple of compensation schedules.

BIBLIOGRAPHY

Baron, D. P. and R. DeBondt. "On the Design of Regulatory Price Adjustment Mechanisms." Journal of Economic Theory. Vol. 24 (1981), pp. 70-94.

Baron, D. P. and R. B. Myerson. "Regulating a Monopolist with Unknown Costs." Econometrica, Forthcoming, (1982).

Baumol, W. J. and Bradford, D. P. "Optimal Departures from Marginal Cost Pricing." The American Economic Review. Vol. 60 (1970), pp. 265-283.

Blackwell, D. "Comparison of Experiments." in Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability. ed. by J. Neyman. University of California Press, 1951.

Harris, M. and Raviv, A. "Optimal Incentive Contracts with Imperfect Information." The Journal of Economic Theory. Vol. 20 (1979), pp. 231-259.

Harris, M. and Townsend, R. M. "Resource Allocation Under Asymmetric Information." Econometrica. Vol. 49 (1981), pp. 33-64.

Holmstrom, B. "Moral Hazard and Observability." The Bell Journal of Economics. Vol. 10 (1979), pp. 74-91.

Loeb, M. and W. A. Magat. "A Decentralized Method of Utility Regulation." The Journal of Law and Economics. Vol. 22 (1979), pp. 399-404.

Sappington, D. "Limited Liability Contracts Between Principal and Agent." Journal of Economic Theory, Forthcoming (1982).

Sappington, D. "Optimal Regulation of a Multiproduct Monopolist With Unknown Costs." University of Michigan Working Paper, December 1981.

Sappington, D. "Strategic Firm Behavior Under a Dynamic Regulatory Adjustment Process." The Bell Journal of Economics. Vol. 11 (1980), pp. 360-372.

Shavell, S. "Risk Sharing and Incentives in the Principal and Agent Relationship." The Bell Journal of Economics. Vol. 10 (1979), pp. 55-73.

Vogelsang, I. and Finsinger, J. "A Regulatory Adjustment Process for Optimal Pricing by Multiproduct Monopoly Firms." The Bell Journal of Economics. Vol. 10 (1979), pp. 157-171.

Weitzman, M. L. "Efficient Incentive Contracts." The Quarterly Journal of Economics. Vol. 94 (1980), pp. 719-730.

CREST Working Papers

C-1 The Ergodic Behavior of Stochastic Processes of Economic Equilibrium, by Lawrence E. Blume.

C-2 The Token Economy, Reinforcement and the Consumer Model, by John G. Cross.

C-3 Notes on a Theory of Learning to Search, by John G. Cross.

C-4 Redistributive Taxation as Social Insurance, by Hal R. Varian.

C-5 Catastrophe Theory and the Business Cycle, by Hal R. Varian.

C-6 Household Bequests, Perfect Expectations, and the National Distribution of Wealth, by John P. Laitner.

C-7 Inflationary Disequilibrium, by John G. Cross.

C-8 Ellet's Transportation Model of an Economy with Differentiated Commodities and Consumers, I: Generic Cumulative Demand Functions, by Carl P. Simon.

C-9 Cournot Equilibrium in Factor Markets, by Theodore Bergstrom.

C-10 Is a Man's Life Worth More Than His Human Capital, by Theodore Bergstrom.

C-11 A Model of Sales, by Hal R. Varian.

C-12 New Techniques for the Study of Stochastic Equilibrium Processes, by Lawrence E. Blume.

C-13 Consistent Expectations, by Lawrence E. Blume.

C-14 On Baubles and Bubbles, by John G. Cross.

C-15 When Does Majority Rule Supply Public Goods Efficiently?, by Theodore Bergstrom.

C-16 "Rational" Duopoly Equilibria, by John P. Laitner.

C-17 The Natural Rate of Unemployment, by John P. Laitner.

C-18 Learning to be Rational, by Lawrence E. Blume and David Easley.

C-19 Notes on Cost-Benefit Analysis, by Hal R. Varian.

C-20 On Capturing Oil Rents with a National Excise Tax, by Theodore Bergstrom.

C-21 Cournot, Novshek and the Many Firms, by Theodore Bergstrom.

C-22 The Nonparametric Approach to Demand Analysis, by Hal R. Varian.

C-23 National Debt, Social Security, and Bequests*, by John P. Laitner.

C-46 Oligopoly Behavior When Conjectural Variations are Rational, by John Laitner.

C-47 Soldiers of Fortune? by Ted Bergstrom.

C-48 Scalar and Vector Maximization: Calculus Techniques with Economics Applications, by Carl P. Simon.