Optimal Rewards for Economic Regulation

By MARTIN L. WEITZMAN*

Suppose several production units or firms must be regulated when costs and benefits are uncertain. Pollution might be a specific example, although there are many others. Given that firms must bear their own costs, the regulators want to transmit a schedule of revenues to each unit which in some expected value sense elicits an optimal response.

What makes this problem intriguing is that while benefits are typically a nonseparable function of *all* the firms' outputs, it seems realistic to require that the revenue function to be received by a given unit must depend in some well-defined way on *its* individual actions alone.

Two control modes often used in regulation are "prices" and "quantitites." These can be viewed as special cases of revenue functions. Prices are a linear function of output. Quantities might be described as a quadratic loss function of deviations from target, accompanied by a heavy-penalty weight. Although these two control modes are frequently treated as mutually exclusive regulatory strategies, it is highly unlikely that either extreme is optimal.

In the class of all objective functions, what is the best revenue schedule? This paper is devoted to formalizing the question, giving a precise answer (at least for an important special case), and analyzing the answer. Roughly speaking, in an optimal policy the center transmits to each firm a "price term" plus a weighted "quantity term," the weight depending in a well-defined way on specific features of the underlying situation. Such a result can be interpreted as providing a reasonable justification for regulation based on *both* price incentives and quantity targets.

*Massachusetts Institute of Technology. On the occasion of his forthcoming 65th birthday, I would like to dedicate this paper to my friend, colleague, and teacher Evsey D. Domar. He fostered my interest in the problem analyzed here by puzzling aloud over the simultaneous presence of price and quantity directives in most planned systems. For their helpful comments, my thanks go to P. A. Diamond, M. Manove, J. M. Mirrlees, and the referee.

I. The Regulatory Environment

As perhaps befits a theoretical paper, "regulation" is being analyzed at a rather high level of abstraction. The basic question is how to make simple rules which will induce firms to do what is best in an uncertain world. This issue is taken as the prototype problem of regulation, and it is modeled below.

The question why an economic activity must be regulated instead of being left to allocate itself in the market place is not treated directly. Possible reasons might range all the way from administrative or political considerations to one form or another of market failure. Prime examples of the kind of regulatory situation I have in mind are control of interdependent divisions in a large organization, and government regulation of externalities. In such situations there is no natural market for the good, and its production must be artificially controlled.

Suppose there are n firms or divisions to be regulated. Let x_i units of commodity i be produced by firm or division i. In the context of an externality, say pollution, x_i would be the level of the ith polluter's abatement program holding everything else constant. Depending on the interpretation, the various components of $x = (x_1, \ldots, x_n)$ might represent physically distinct goods or they could denote amounts of the same item produced by different production units.

The word "commodity" is being used in an abstract sense and really could pertain to just about any kind of good from pure water to military hardware. For the sake of preserving a unified notation we follow the standard convention of treating goods as desirable. Rather than talking about air pollution, for example, I instead deal with its negative—clean air.

For a firm to produce output requires the outlay of a corresponding cost. An essential feature of the regulatory environment I am trying to describe is uncertainty about the exact specification of each firm's cost function. In most cases even the managers and engineers most closely associated with production would be unable to precisely specify beforehand the cheapest way of generating various hypothetical output levels. Because they are yet further removed from the production process, the regulators are likely to be vaguer still about a firm's cost function. This observation acquires additional force in a fast moving world where deception may be involved or where knowledge of particular circumstances of time and place may be required.

Generally speaking, there is no way the regulators can know beforehand exactly what it will cost to achieve a certain output level. Estimates can be made and the degree of fuzziness could be reduced by investigation and research. But it could never be eliminated completely because new sources of uncertainty are arising all the time. The true costs will only be known when production is actually underway.

In mathematical language, the regulators perceive the cost function of firm i as an estimate or approximation, written

$$C_i(x_i; \epsilon_i) \tag{1}$$

In the above formulation ϵ_i is a disturbance term, stochastic element, or random variable representing a state of the world unobserved and unknown at the present time. During the course of plan implementation, ϵ_i will eventually make itself known to firm i, and perhaps also to the regulators. But at the moment when an operational plan must be decided for the forthcoming period, the regulators' knowledge of ϵ_i can be represented only by a probability distribution.

The benefit function too is presumably discernable only tolerably well, say as

$$B(x; \delta) \tag{2}$$

with δ a vector of random variables having some probability distribution. The money value of various commodity output levels may be uncertain because it is imperfectly known or because authentic randomness (like the weather) is present.

It is assumed that C_i is strictly convex in x_i for each ϵ_i and B is strictly concave in x for each δ. All cost and benefit functions are presumed to be smoothly differentiable.

II. A Problem in Regulation

There is another important feature of cost functions that goes along with the uncertainty. Not only are costs unknown, but it is typically difficult and expensive to find out what they are. Sometimes economists and others share an overtendency to conceptualize regulation as a process of continual fine tuning. A certain strategy is adopted, then marginal costs and marginal benefits are observed. If they are not equal, the fees, standards, or other parameters are smoothly adjusted until an optimum is obtained.

However, this is an inappropriate way of viewing the problem. In order to be given a chance to work, a regulatory strategy must be left in place for an extended period after it has been adopted. If a firm anticipates the regulations are going to change in the near future, it is not going to take very seriously compliance with them now. This does not mean that regulations, once formulated, must be immutable for all time. It is just that they must remain in force long enough to be believable.

Another, perhaps more serious, reason that the fine tuning model may be irrelevant is that most production activity involves investment. The investment may be in research, development, reorganization, new equipment, learning by doing, etc. True costs will not become known until the investments are actually made. Whatever its form, such investment takes time and it is largely irreversible. Once made, it cannot be easily or costlessly taken back, nor can the knowledge gained be effortlessly transferred to other situations. This means that there are costs to adjusting regulations, and they are likely to be substantial.

A basic principle of regulation is that the regulators are forced to make decisions in an uncertain environment and they must live with the consequences for some time. Among these consequences is the possibility that costs borne by some firms will turn out to be higher or lower than was expected. A good regulatory strategy will take advantage of this by instituting a reward structure which automatically encourages the cheap firm to produce more and the expensive firm less. In our formulation, regulators are confined to a strategy of indirect control by judiciously selecting revenue functions in advance for each firm.

Now, in a certain sense the ideal revenue function for any firm is the entire expected benefits function, plus or minus some constant. Assuming away the game-theoretic problems having to do with bluffing, threatening, etc., a Nash-type equilibrium might conceivably emerge where each firm would have the incentive to set its marginal cost equal to its marginal benefit after all uncertainty had been eliminated and every firm knew what every other firm was doing.

The trouble with this sort of approach is that benefits are typically a nonseparable function of *all* the firms' outputs, whereas a particular firm has control only over its *own* output. It seems like a relevant abstraction to insist that a regulatory agency cannot reward or penalize a firm in what might be viewed as an arbitrary or capricious manner. Asking a firm to bear the extra risk involved in adopting a revenue schedule depending on uncertain variables not under its control may be infeasible or unacceptable. Some of the reasons for this have just been cited in downplaying the relevance of the fine tuning model. In addition, such a schedule may simply be too complicated to handle.

That revenue functions should depend only upon individual actions is a strong assumption (for example, it rules out profit-sharing incentive schemes), but I think it is appropriate to the kind of regulatory environment I have described. In this paper I take as a point of departure a scenario where firms pay their own costs and the state sets revenue functions for each firm which depend *only* on that firm's output.

III. A Formulation of the Basic Problem

A *revenue function* $R_i(x_i)$ is a schedule of monetary payments received by firm i as a function of its output. For example, if a price p_i is paid for the output of firm i, the corresponding revenue function is

$$R_i(x_i) = p_i x_i \tag{3}$$

Or, if it is the intention of the planners to set a quota $\hat{x}_i$, they might specify the following revenue function:

$$R_i(x_i) = \frac{-q_i}{2}(x_i - \hat{x}_i)^2 \tag{4}$$

where q_i is a large number.

It is important to realize that the process of profit maximization causes every revenue function to generate some output response. The response depends on the revenue function $R_i(\ )$ and the state of the world ϵ_i. For a given $R_i(\ )$ and ϵ_i, firm i will set its output x_i at that level which maximizes profits, implicitly solving the equation

$$\max_{x_i \geq 0} R_i(x_i) - C_i(x_i; \epsilon_i) \tag{5}$$

Equation (5) should not be interpreted too literally as saying that the firm knows the exact value of ϵ_i with certainty (at the same time the regulator knows only the probability distribution of ϵ_i). In the scenario I have in mind, when a revenue function is instituted for a sufficiently long period the firm will eventually grope its way to a profit-maximizing output, presumably by trial and error testing of the relevant alternatives. This is quite a different interpretation from having the cost function known a priori.

Without any significant loss of generality in the problem to be posed, we limit attention to revenue functions which generate *unique* output responses. That is, the solution of (5) is some response *function*

$$x_i = G_i(R_i(\ ), \epsilon_i) \tag{6}$$

satisfying for all possible ϵ_i the condition

$$R_i(G_i(R_i(\ ), \epsilon_i)) - C_i(G_i(R_i(\ ), \epsilon_i); \epsilon_i) = \max_{x_i \geq 0} R_i(x_i) - C_i(x_i; \epsilon_i) \tag{7}$$

Note that changing a revenue function by adding or subtracting any constant cash payment does not alter the corresponding response (aside from the issue of setting such a low payment that the firm is forced out of business altogether). At least in a rough way, this might be interpreted as providing some justification for studying the allocative effects of a revenue function apart from the distributive consequences.

In the framework adopted here, the planners are at a point where as much information as is feasible to gather has already been obtained. An operational plan must now be decided on the basis of the available current knowledge, summarized by (1) and (2). Because it will force long-term resource commitments (like capital investments), any incentive scheme has serious consequences which continue for some time and cannot easily be reversed. This is an essential feature of the regulatory environment very prominent, for example, in the case of pollution. A regulatory agency must resign itself to naming in advance revenue functions $\{R_i(\quad)\}$ and living with the outcome even though it does not presently know the values of $\{\epsilon_i\}$ or δ.

Through the output response (6) which they induce, reward functions $\{R_i(\quad)\}$ yield the expected differences in benefits and costs

$$(8)\quad \Phi(\{R_i(\quad)\}) \equiv \underset{\{\epsilon_i\},\delta}{E}\Big[B(\{G_i(R_i(\quad),\epsilon_i)\};\delta) - \sum_{i=1}^{n} C_i(G_i(R_i(\quad),\epsilon_i);\epsilon_i)\Big]$$

A set of *optimal* revenues $\{R_i^*(\quad)\}$ is any collection of functions which maximize (8). In other words, via the output response generated by them, optimal revenue functions maximize expected benefits minus costs. This can formally be written[1]

$$(9)\quad \Phi(\{R_i^*(\quad)\}) = \max_{\{R_i(\quad)\}} \Phi(\{R_i(\quad)\})$$

[1]The maximization is over the class of all possible reward functions yielding response functions. It is not difficult to prescribe conditions which ensure the existence of a solution to (9).

The above problem shares certain features of the more general structure analyzed in the theory of teams. Indeed, one of the more significant results of team theory will be used in proving the basic theorem of this paper.

Note that (9) is easy to solve when the benefit function is additively separable in the output of each firm. Then the optimal revenue function for a firm is just *its* part of expected benefits. The interesting case is where the benefit function is not separable.

IV. Optimal Revenue Functions

The remainder of this paper is devoted to characterizing the form of an optimal revenue function and explaining its dependence on various factors. Under the most general circumstances this appears to be a very intricate task. Fortunately a complete characterization is possible for an important special case.

Optimal revenue functions generate a range of output responses as the uncertainty varies. From now on it will be assumed that within this output range marginal costs and marginal benefits can be accurately approximated by linear forms.

A linear approximation might be rationalized on one of two grounds. The amount of uncertainty could be small enough to keep the range of output responses sufficiently limited to justify a first-order approximation. Or, it might just happen that total cost and benefit functions are almost quadratic to begin with. At any rate, the possibility of sharply characterizing an optimal solution makes the linear case a natural preliminary to any more general analysis.

Consider for a moment the problem of finding an optimal set of quotas or targets $\hat{x} = (\hat{x}_1, \ldots, \hat{x}_n)$. The optimal quota maximizes expected benefits minus expected costs, so that

$$(10)\quad \underset{\{\epsilon_i\},\delta}{E}\left[B(\hat{x},\delta) - \sum_{i=1}^{n} C_i(\hat{x}_i;\epsilon_i)\right] = \max_{x} \underset{\{\epsilon_i\},\delta}{E}\left[B(x;\delta) - \sum_{i=1}^{n} C_i(x_i;\epsilon_i)\right]$$

Presuming it is interior, the solution of (10) satisfies the first-order condition

$$(11)\quad p_i \equiv \underset{\delta}{E}\, B^i(\hat{x};\delta) = \underset{\epsilon_i}{E}\, C_i'(\hat{x}_i;\epsilon_i) \qquad i = 1,\ldots,n$$

where p_i is the expected marginal benefit equals marginal cost of the ith commodity evaluated at the optimal quota.[2]

Under the linearity assumption, the marginal cost of the ith producer can be written

$$(12)\quad C_i'(x_i;\epsilon_i) = p_i + \gamma_i(x_i - \hat{x}_i) + \epsilon_i \qquad i = 1,\ldots,n$$

while the marginal benefit of commodity i can be expressed as

$$(13)\quad B^i(x;\delta) = p_i - \sum_{j=1}^{n} \beta_{ji}(x_j - \hat{x}_j) + \delta_i \qquad i = 1,\ldots,n$$

where, without loss of generality,

$$(14)\quad E\delta_i = E\epsilon_i = 0 \qquad i = 1,\ldots,n$$

The various δ_i are just components of δ.

In order to obtain sharp results, a further regularity assumption on the probability distributions is needed. The conditional expectation of ϵ_j given ϵ_i is presumed proportional to ϵ_i. Likewise for the expected value of δ_i conditional on ϵ_i. That is,

$$(15)\quad \begin{aligned} E\epsilon_j/\epsilon_i &= \theta_{ji}\epsilon_i \qquad i = 1,\ldots,n \\ E\delta_i/\epsilon_i &= \eta_i\epsilon_i \qquad j = 1,\ldots,n \end{aligned}$$

for some coefficients $\{\theta_{ji}\}$, $\{\eta_i\}$. Naturally $\theta_{ii} = 1$.

Condition (15) could be justified as a first-order approximation holding for small uncertainties. It would also be a consequence of a joint normal distribution in $\{\epsilon_i\}$ and $\{\delta_i\}$. All independent probability distributions ($\theta_{ji} = 0$, $j \neq i$) automatically satisfy (15).

Note that

$$\theta_{ji} = \sigma_{ij}^2/\sigma_{ii}^2,\ \eta_i = \sigma_{i0}^2/\sigma_{ii}^2$$

where $\quad \sigma_{ij}^2 = E\epsilon_i\epsilon_j,\ \sigma_{i0}^2 = E\epsilon_i\delta_i$

[2]All the $\{p_i\}$ are identical when the various commodities represent the same item produced by different production units.

The basic result of the present paper is summarized by the following:

THEOREM: *Under the assumptions* (12)-(15), *the optimal reward function for unit i can be expressed in the form*

$$(16)\quad R_i^*(x_i) = p_i x_i - \frac{q_i}{2}(x_i - \hat{x}_i)^2 \pm constant$$

The $\{q_i\}$ *are coefficients satisfying*[3] *the equations (linear in* $\{1/q_i + \gamma_i\}$*)*

$$(17)\quad \sum_{j=1}^{n} \frac{\beta_{ji}\theta_{ji}}{q_j + \gamma_j} + \eta_i = 1 - \frac{\gamma_i}{q_i + \gamma_i} \qquad i = 1,\ldots,n$$

V. Analysis of an Optimal Reward

Equation (16) means that aside from the arbitrary constant, an optimal reward function can be decomposed into two components.

The first term

$$(18)\qquad p_i x_i$$

is the traditional price signal. If p_i accurately represented the marginal benefit of commodity i, using (18) as a reward function would automatically induce firm i to produce at that output level where marginal benefit equals marginal cost. The apparent guarantee of social efficiency is what makes the use of prices as a regulatory device so attractive to the economist. Unfortunately for this idea, the marginal benefit of commodity i cannot be reduced to a single number which is precisely known beforehand.

[3]It is assumed that (17) has a solution and that $\{q_i + \gamma_i\}$ are positive. The latter condition is needed to guarantee that the problem of maximizing revenues (16) minus costs has a meaningful solution for each firm. Although it seems hard to prove a very broad sufficiency theorem, playing with lots of examples has convinced me that the condition holds for most cases of economic interest. A sufficiency theorem can be proved when firms are close to being symmetric with each other or as the uncertainties are not too far from being independently distributed.

The second component of (16),

$$\frac{-q_i}{2}(x_i - \hat{x}_i)^2 \tag{19}$$

is a quadratic penalty for departures from the target $\hat{x}_i$. Were $\hat{x}_i$ in fact the socially optimal output of commodity i, the center could do no better than transmitting (19) as a revenue function, with q_i arbitrarily large (which is equivalent to setting $\hat{x}_i$ as a standard). The seeming ability to directly fix economic activity at the socially desirable level is what makes the quota appealing as a regulatory device, especially to the general public. Alas, the regulators don't know exactly what output levels are socially optimal to begin with.

The basic result of this paper argues that in economic planning situations prices and quotas are not redundant or inconsistent messages. In fact, a "mixed" price-quota system is the optimal reward. The coefficient q_i determines the composition of the mix. With $q_i = 0$, (16) becomes a pure price signal; when $q_i = \infty$, (16) is made into a complete quota system.[4]

All this strongly suggests that a regulatory strategy based on *both* price incentives and quantity targets, far from being a contradiction, is actually optimal in a world of uncertainty. Such a principle has been intuitively sensed, I believe, by practical planners. Of course the revenue function is usually not formalized as it is in (16). Instead there is typically some vaguely ambiguous policy of rewarding output while simultaneously discouraging deviations from a target. Taking advantage of the theorist's inherent right of simplification, I would suggest that (16) is not a bad translation of such a policy.

A result like (16) provides at least a partial resolution of the environmental economics debate between the use of effluent charges and the use of effluent standards. It also offers some justification for mixed price and quantity controls within a large, divisionalized organization.

[4]The comparative advantage of these two extreme regulatory modes was analyzed in my 1974 article.

Note that the optimal reward function does not promise social optimality or efficiency *ex post*, after $\{\epsilon_i\}$ and $\{\delta_i\}$ take on specific values. The concept of *ex post* social optimality is too strong to require, given the informational constraints being imposed. The relevant issue is which reward function comes closest to inducing a social optimum in some average sense.

A simple description explains how $\{\hat{x}_i\}$ and $\{p_i\}$ are determined. They are just the optimal outputs and their marginal values obtained when the center, suppressing all uncertainty, maximizes the difference between the "representative" benefit function

$$b(x) \equiv \underset{\delta}{E}\ B(x;\delta)$$

and the sum of "representative" cost functions

$$c_i(x_i) \equiv \underset{\epsilon_i}{E}\ C_i(x_i;\epsilon_i)$$

The determination of the penalty weights $\{q_i\}$ is slightly more complicated to explain. Differentiating (16) and setting the resulting expression equal to (12) yields the response function

$$x_i(\epsilon_i) = \hat{x}_i - \frac{\epsilon_i}{q_i + \gamma_i} \tag{20}$$

Suppose $\epsilon_i = 1$. This lowers the output of x_i by $1/(q_i + \gamma_i)$ units, causing a net increase of

$$1 - \frac{\gamma_i}{q_i + \gamma_i} \tag{21}$$

dollars in the marginal cost of firm i. By (15), δ_i is expected to be η_i dollars above average while ϵ_j is expected to be θ_{ji} dollars above its mean. From (15) and (20), firm j can be expected to curtail output by $\theta_{ji}/(q_j + \gamma_j)$ units. Using (13), the marginal benefit of commodity i is expected to increase by

$$\eta_i + \sum_{j=1}^{n} \frac{\beta_{ji}\theta_{ji}}{q_j + \gamma_j} \tag{22}$$

dollars. Equating the increase in marginal costs (21) with the expected increase in marginal benefits (22) yields condition (17).

On what things do the coefficients $\{q_i\}$ depend? To strengthen our intuitive feeling for the meaning of equation (17), let us turn first to a special case which can serve as a point of departure.

Suppost that all the uncertainties are independent, so that

$$(23) \qquad \theta_{ji} = 0 \qquad \text{for } j \neq i$$
$$\eta_i = 0$$

In this case equation (17) reduces to

$$(24) \qquad q_i = \beta_{ii}$$

Under (23) the optimal penalty coefficient for a commodity is just the curvature of the benefit function in that commodity. With independent uncertainties firm i should be given as a reward that part of expected benefits which remains as a function of *its* output when the outputs of all other firms $j (j \neq i)$ have been parametrically fixed at $\hat{x}_j$. This interpretation comes from examining (13), (16), and (24).

The greater the curvature in benefits, the more significant is the weight of the quantity term (19) in the reward function mix (16). If marginal benefits decrease rapidly around the optimal quota, there is a high degree of risk aversion and the center cannot afford being even slightly off the mark. Relying too much on the price mode is risky because a miscalculation results in under or overshooting the target, with detrimental consequences. In such a situation the quantity mode scores a lot of points because a high premium is put on the rigid output controllability which only it can provide under uncertainty.

On the other hand, the weight of the quantity term is lessened when benefits are closer to being linear. In that case it would be foolish to place too much emphasis on targets. Since expected marginal social benefit is approximately constant over some range, a superior policy comes closer to naming it as a price and letting the producer find the optimal output level himself after eliminating the uncertainty from costs.

Returning to the more general case, when (23) does not hold q_i will tend to exceed β_{ii}. Generally speaking, the penalty coefficients $\{q_i\}$ become more significant as the interdependence coefficients $\{\theta_{ij}\}$ and $\{\eta_i\}$ increase.

The reason for this is easy to understand. A positive θ_{ji} means that when the marginal costs of firm i are low, so are those of firm j. Whenever firm i is increasing output because its costs are low, firm j is doing likewise. Compared with a situation of independent marginal costs, more damping should be introduced; this would stabilize welfare decreasing over and underreactions in aggregate output responses.

Something analogous happens in the case of positive η_i. Producers will cut back output for higher marginal costs, but this cutback should be dampened when there tends to be a simultaneous increase in marginal benefits. In such situations a greater weight for the quantity mode is appropriate because that mode has better properties as a stabilizer. The story is the other way round when η_i is negative.

The coefficient β_{ij} is a measure of the degree of complementarity between commodities i and j. When it is higher, more stabilizing is desirable to keep commodities i and j closer to their appropriate proportions. When it is lower, less weight is needed on individual quantity terms because greater substitutability is possible.

An instructive illustration of what the quantity weights depend on is provided by the special regularized case of perfect symmetry:

$$\gamma_i = \gamma \qquad \eta_i = \eta$$
$$\theta_{ii} = 1 \qquad \beta_{ii} = \beta$$
$$\theta_{ij} = \rho \quad i \neq j \qquad \beta_{ij} = \mu\beta \quad i \neq j$$

In this case (17) yields

$$q_i = q = \frac{\eta\gamma + \beta + (n-1)\mu\beta\rho}{1-\eta}$$

The quantity weight q increases in β, ρ, η, and μ, verifying our previous discussions.

VI. Proof of the Main Proposition

Consider the problem of finding a set of optimal response functions $\{x_i(\epsilon_i)\}$ in an information structure where firm i observes only ϵ_i and controls only x_i. The "objective function" is the expected difference between benefits and costs. The problem is to maximize over $\{x_i(\epsilon_i)\}$ the function

$$(25)\quad \psi(\{x_i(\epsilon_i)\}) = \underset{\{\epsilon_i\},\delta}{E}\left[B(\{x_i(\epsilon_i)\};\delta) - \sum_{i=1}^{n} C_i(x_i(\epsilon_i);\epsilon_i)\right]$$

Since by assumption (6) any revenue function generates a response function, the solution to the above problem yields at least as high a value of the objective function as the solution to problem (9).

Now it turns out that finding optimal response functions in the present framework is an example of a classical problem in the theory of teams, whose solution[5] is given by (20) with the definition (17). Because the general result is typically presented in a somewhat different framework and may be difficult to follow, I will sketch a proof for the case treated here.

For simplicity, assume discrete distributions. Let ϵ_{it} be a value which ϵ_i takes on with positive probability. Let $x_{it} = x_i(\epsilon_{it})$ be the output response of firm i when $\epsilon_i = \epsilon_{it}$. It is not difficult to show that $\psi(\ )$ in (25) is a concave differentiable function of the variables $\{x_{it}\}$ over which it is being maximized (this derives essentially from the concavity-differentiability of the benefit minus cost function and the concavity-differentiability preserving properties of an expected value operator). Hence the appropriate first-order conditions are necessary and sufficient for an optimum.

[5]See Jacob Marshak and Roy Radner, Theorem 5, p. 168. Matching up my notation with theirs is a bit messy, and it seemed better to omit the details, which the interested reader should be able to supply. I am indebted to Kenneth J. Arrow for pointing out to me that my characterization of optimal response functions is really a special case of Radner's result.

From (12), the marginal expected cost of x_{it} given $\epsilon_i = \epsilon_{it}$ is

$$(26)\quad E\left[\frac{\partial C_i}{\partial x_{it}}\middle|\epsilon_{it}\right] = p_i + \gamma_i(x_{it} - \hat{x}_i) + \epsilon_{it}$$

From (13), the marginal expected benefit of x_{it} given $\epsilon_i = \epsilon_{it}$ is

$$(27)\quad E\left[\frac{\partial B}{\partial x_{it}}\middle|\epsilon_{it}\right] = p_i - \sum_{j=1}^{n} \beta_{ji}(E[x_j \mid \epsilon_{it}] - \hat{x}_j) + E[\delta_i \mid \epsilon_{it}]$$

Equating (27) with (26) yields the first-order condition

$$(28)\quad -\sum_{j=1}^{n} \beta_{ji}(E[x_j \mid \epsilon_{it}] - \hat{x}_j) + E[\delta_i \mid \epsilon_{it}] = \gamma_i(x_{it} - \hat{x}_i) + \epsilon_{it}$$

We must verify that the solution proposed in (20) satisfies (28). Therefore, for x_j substitute

$$(29)\quad x_j = \hat{x}_j - \frac{\epsilon_j}{q_j + \gamma_j}$$

and for x_{it} substitute

$$(30)\quad x_{it} = \hat{x}_i - \frac{\epsilon_{it}}{q_i + \gamma_i}$$

Plugging (29) and (30) into (28) yields

$$\sum_{j=1}^{n} \frac{\beta_{ji} E[\epsilon_j \mid \epsilon_{it}]}{q_j + \gamma_j} + E[\delta_i \mid \epsilon_{it}] = -\frac{\gamma_i \epsilon_{it}}{q_i + \gamma_i} + \epsilon_{it}$$

Using (15), the above expression becomes

$$(31)\quad \epsilon_{it} \sum_{j=1}^{n}\left(\frac{\beta_{ji}\theta_{ji}}{q_j + \gamma_j} + \eta_i\right) = \epsilon_{it}\left(1 - \frac{\gamma_i}{q_i + \gamma_i}\right)$$

Equation (31) will hold for all possible ϵ_{it} if the $\{q_i\}$ are defined by (17). Thus, expression (20) is indeed the optimal response function.

The remainder of the proof consists of verifying that under cost assumption (12), the revenue function (16) generates[6] the response function (20).

[6]Recall I am assuming the $\{q_i + \gamma_i\}$ satisfying (17) to be positive. If the solution to (17) yields a negative value of $q_i + \gamma_i$ for some i, there will still exist an optimal response function (in a team theory sense), given by (20). Unfortunately, there is no way to induce firm i to follow this rule by naming a corresponding revenue function. The problem is that with a negative $q_i + \gamma_i$, (20) dictates that the firm should produce *more* when its costs are *higher*. This kind of seemingly perverse behavior might be optimal if, for example, whenever the costs of firm i are high, the costs of other firms are much higher still. But no revenue function can elicit such a perverse response from firm i. The regulators would have to rely on moral suasion or some other means. The team theory approach would make no distinction between positive and negative $q_i + \gamma_i$. But the reliance on revenue functions to elicit proper behavior requires that $q_i + \gamma_i$ be positive for all firms, at least for the theorem proved here.

REFERENCES

Jacob Marshak and Roy Radner, *Economic Theory of Teams*, New Haven 1972.

→ **M. L. Weitzman,** "Prices vs. Quantities," *Rev. Econ. Stud.*, Oct. 1974, *41*, 477–91.