ISSN 1045-6333

INNOVATION UNDER TECHNOLOGY-BASED REGULATION

Bruce L. Hay

Discussion Paper No. 168

9/95

Harvard Law School
Cambridge, MA 02138

The Program in Law and Economics is supported by
a grant from the John M. Olin Foundation.

INNOVATION UNDER
TECHNOLOGY-BASED REGULATION

Bruce L. Hay
Harvard Law School

Abstract

Regulatory standards are frequently based on firms' ability to control externalities, for example their ability to control pollution emissions. In such systems, the regulated industry may have little to gain by improving its control technology, because the regulator will respond by imposing more rigorous standards. The paper examines the incentives of firms to adopt efficient control technologies in this setting. It shows how the design of regulatory standards -- in particular, the way in which firms with relatively efficient control technologies are treated vis-a-vis other firms -- influences firms' choice of technology. It attempts to identify the circumstances in which firms may rationally choose to improve their control technology even when doing so is not in their collective interest.

INNOVATION UNDER TECHNOLOGY-BASED REGULATION

Bruce L. Hay*

I. INTRODUCTION

Regulation of externalities typically consists of standards that are based on firms' ability to take precautions. For example, pollution control laws require firms to reduce pollutant emissions to the extent permitted by "available technology." Occupational safety laws require firms to control workplace risks to the extent "feasible." Negligence standards in tort law require firms to take cost-justified measures against harm. In these examples and others, rules stating what firms *must* do rest on regulators' perception of what they *can* do.

This approach to regulation, sometimes termed "technology-based," contains a built-in problem of moral hazard. Generally speaking, society benefits if industry improves its technology for reducing hazards. But innovation of this type may leave industry worse off, because the improvement in control technology simply leads to greater regulation. And there is often no way for the regulator to observe (and penalize) "forgone" innovation.

The purpose of this paper is to examine the contours of this problem. The question I address is: Under what circumstances will firms rationally undertake innovation of this kind, given that the regulatory authority will respond by tightening the regulatory standard? My objective is to show how the structure of regulation affects firms' incentives to improve their control technologies.

---

*Harvard Law School. I thank Louis Kaplow and Richard Revesz for commenting on an earlier version.

To do this, I use a model in which firms make technological decisions strategically in the face of the anticipated reaction of the regulator. The model examines firm behavior under three "ideal types" of technology-based standard, intended to capture the esssential features of standards actually encountered in practice. These ideal types are: (i) standards that differentiate among firms by imposing greater control obligations on the more efficient firms; (ii) uniform standards that require all firms to match the optimal performance of the *average* firm in the industry; and (iii) uniform standards that require all firms to match the optimal performance of the *most efficient* firm in the industry.

A central conclusion is that uniform standards tend to encourage firms to improve their control technology even when doing so is not in their collective interest -- and indeed even when doing so is not in society's interest. The essential insight is simple: if other firms refrain improving their technology, a given firm may be tempted to "break ranks" and improve its own technology, for the sake of reducing its costs either in *absolute* terms or *relative to other firms*. As a result, equilibria in which firms forgo technology improvements may be difficult to sustain.

This happens for different reasons under the two types of uniform standard in the model. If the standard is geared to the *average* firm in the industry, the standard is relatively insensitive to any individual firm's choice of technology. Each firm may then seek to lower its costs by improving its technology -- leading, by a process of adverse selection, the entire industry to improve its technology. In effect, each firm's effort to lower its own costs leads to higher costs for all.

If, instead, the standard is geared to the *most efficient* firm(s) in the industry, the standard may be quite sensitive to an individual firm's action. Adopting a technology more efficient than any in use may lead to a substantial tightening of the regulatory standard. Nonetheless, an individual firm may want to adopt such a technology, if the effect is to obtain a cost advantage over its less efficient rivals. Though the firm does not reduce its costs in absolute terms by improving its technology, it does reduce its costs *relative to rival firms*. Competition to obtain a cost advantage (and avoid being put at a competitive disadvantage) may lead all firms to improve their technology.

To explore these points, I model the firms in the industry as choosing between two control technologies, one more efficient than the other. Section II of the paper introduces the model by showing the potential divergence in interest between industry and society as regards industry's choice of technology. Section III describes the three ideal types of standard as they operate in the model, while Section IV analyzes the different payoffs these standards create for innovation. Section V derives the main results of the paper, and Section VI discusses applications of the results.

To clarify my subject matter, I should emphasize that the analysis focuses on different types of technology-based standard; it does *not* examine alternatives to technology-based regulation, such as taxes, marketable permits, and so on. The innovation problems I discuss in the paper would, in theory, largely be solved by implementing such a market-based alternative to conventional regulation. But technology-based standards are sufficiently pervasive that it is worth analyzing them in their different forms, putting to one side the (widely discussed) question of how alternative approaches to regulation might do a better job.

II. THE PROBLEM: MORAL HAZARD AND TECHNOLOGY-BASED REGULATION

Consider a simple model in which a government agency is charged with regulating some harmful externality, say, industrial pollution. The agency's objective is to promulgate the standard that maximizes social welfare, defined here as the social benefits of improved air quality net of the industry's costs of controlling its emissions. Regulation takes the form of a standard that limits each factory's emissions to some specified amount (so many pounds of pollution per unit of production, for example). We assume the industry is composed of homogeneous firms.

More than one technology may be available to the industry for removing pollution from its emissions. It might, for example, use cleaner production inputs; it could modify its production process to generate less waste; it might install some machinery in its smokestacks to eliminate the pollutants. These technologies vary in their cost and effectiveness in controlling factory emissions. Let us assume that whatever technology the industry selects becomes the basis for the agency's standard.

In this setting, the industry's interest in using efficient control technologies is ambiguous.[1] Other things equal, the industry is better off using the most efficient available technology — the one that removes each unit of pollution at least cost. But if it does so, then the agency will respond by imposing relatively onerous control obligations on the industry, precisely because its members can reduce its emissions at relatively low cost. In contrast, by using a relatively

---

[1]This has frequently been pointed out in the past, most cogently in the work of Richard Stewart. See, for example, his *Regulation, Innovation, and Administrative Law: A Conceptual Framework*, 69 Calif. L. Rev. 1256 (1981); My intended contribution in this subpart is to identify the formal conditions in which the industry will or will not find innovation attractive.

inefficient control technology, the industry will induce the agency to impose a comparatively lenient standard.

To make the problem concrete, suppose that the industry (all of its members) use some pollution control technology called Alpha. Suppose, in addition, that there exists some new control technology called Beta. This new technology is more efficient than Alpha, in that it reduces the marginal cost of removal. (From now on I will refer to Alpha as the "inefficient" technology and Beta as the "efficient" technology.)

This effect of introducing the new technology is depicted in Figure 1. The horizontal axis represents the quantity of pollution each firm *removes* from its emissions over a given period of production. If the industry uses the inefficient technology, the agency sets a standard (s_α) corresponding to the optimal use of that technology. If the industry were to adopt Beta, the agency would set a new standard (s_β) corresponding to the optimal use of *that* technology. Let us assume that, in addition to the costs depicted in Figure 1, each firm would incur some capital cost k if it switched to the efficient technology.

Using Figure 1, it is easy to see why the industry's interests may diverge from the society's in choosing a technology. For each production period, net social benefits are represented by the area under the marginal benefit curve *minus* the area under the relevant marginal cost curve. Thus, industry's adoption of the efficient technology augments overall social welfare whenever $F+G+H - k$ is greater than H,[2] that is, whenever

$$F + G > k. \qquad \textbf{(1)}$$

---

[2]If the industry keeps Alpha and the agency sets (or keeps) the corresponding standard, then social welfare is given by the area of the triangular region H. If the industry adopts Beta and the agency sets the corresponding standard, then social welfare is given by $F+G+H - k$.

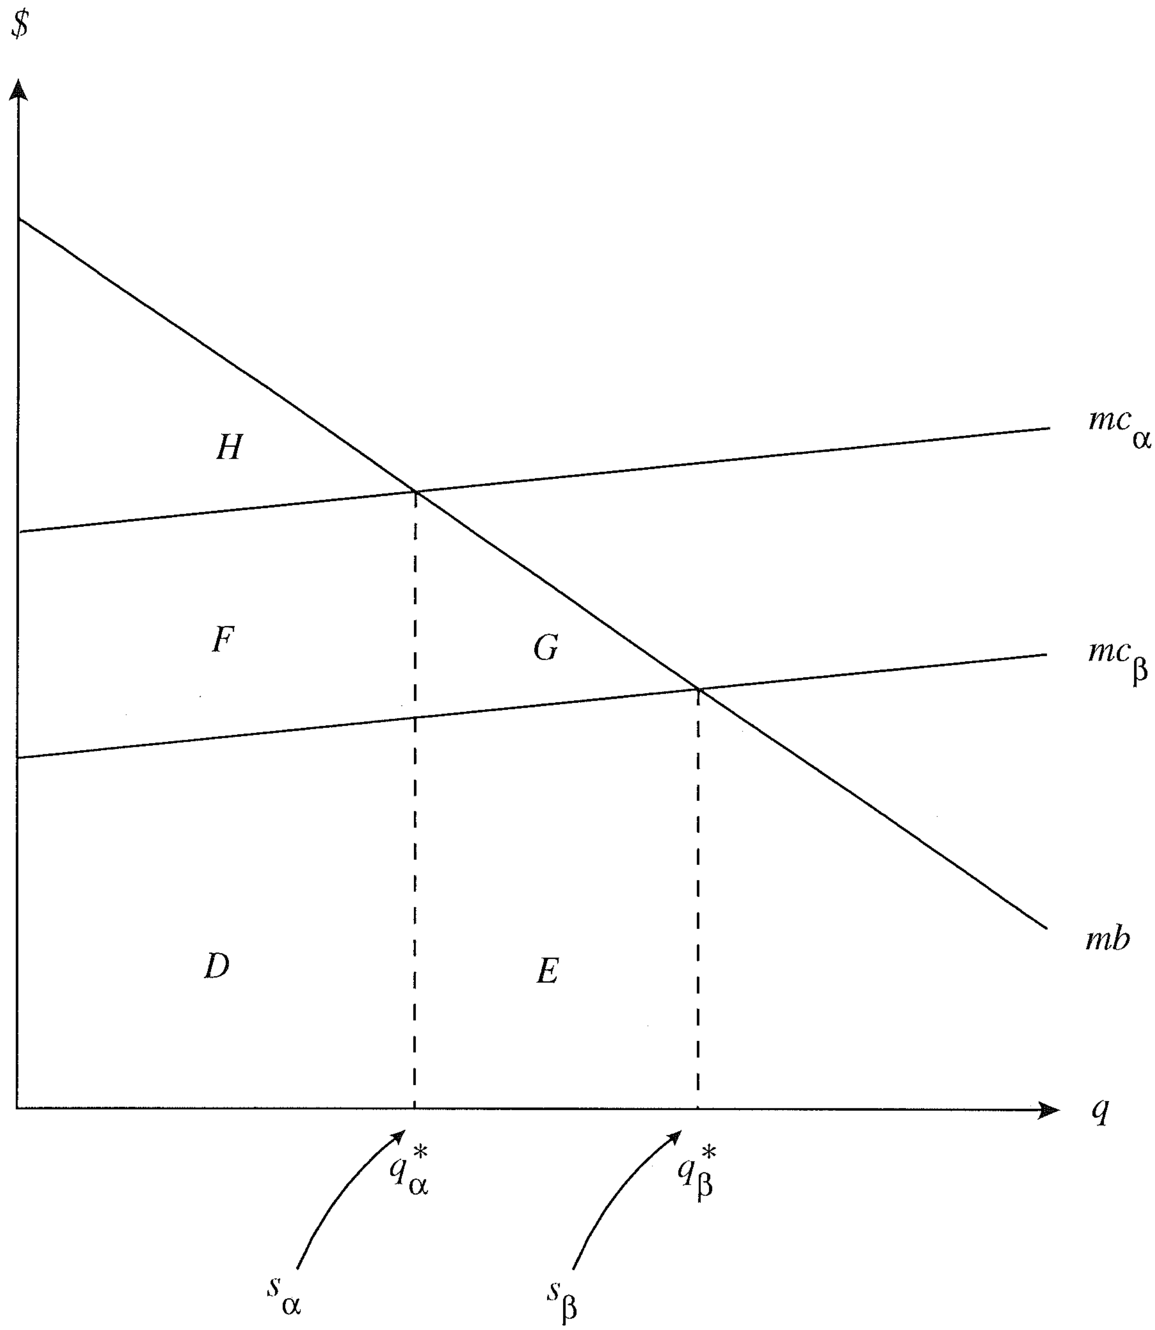

FIGURE 1. – Costs and benefits of alternative control technologies. Notation: q = quantity of emissions removed; mb = marginal social benefit of emissions removal; mc_α = marginal cost of emissions removal using Alpha; mc_β = marginal cost of emissions removal using Beta.

In contrast, the costs to industry are simply given by the area under the relevant marginal cost curve. Adopting the efficient technology lowers industry's control costs only if $D+F$ is greater than $D+E+k$, that is, if

$$F - E > k. \qquad \textbf{(2)}$$

The interests of society and industry diverge, then, whenever **(1)** but not **(2)** is satisfied.

Now, though I use emissions control as the example, this model can be applied to many other regulatory contexts. Alpha and Beta might be thought of as consumer product safety devices; as workplace safety measures; as hazardous waste disposal measures; and so forth. The institutional context may be command-and-control regulation by an agency, tort or contract litigation in court, even commercial or labor arbitration.

In any of these settings, two conditions suffice to give the regulated industry an interest in using the inefficient control technology, even though society would be better off if it used the efficient technology. First, the regulatory standard varies with the technology in use, so that using the efficient technology increase industry's control obligations. Second, the industry does not internalize the social benefits of it added control efforts. Relax either feature, and the problem goes away; but both features are present in many regulatory settings.

Consider the question of internalization. Assume that **(1)** is satisfied but **(2)** is not. In principle, industry could be given a side payment for using the efficient technology, thus enabling the industry to internalize (some of) the social benefits of its control efforts, and so leaving everybody better off.[3] In many settings, however, transaction costs are too high for the beneficiaries of regulation to make this side payment directly,[4] and constraints of administrative or political feasibility prevent the regulatory agency (or some other government authority) from doing so on behalf of the beneficiaries.[5] We will assume in the analysis to follow that side payments or other internalization measures -- including market-based alternatives to regulation, such as taxes or tradeable permits -- are not available.ization internalizatio

What about the first feature -- standards that vary with firms' observed capabilities? Dropping this feature might make it in industry's interest to adopt the efficient technology, since by definition its overall compliance costs would go down. However, making the standard independent of the technology in use is likely to produce misallocations, if the agency is uncertain about what technologies are available to industry. In this model, if the agency guesses wrong about Beta's existence, it risks underregulating (if it wrongly concludes Beta does not exist) or overregulating (if it wrongly concludes Beta does exist).[6] Hence regulators' widespread reliance on observed capabilities as the benchmark for standards.

---

[3]Paying each firm any amount greater than $F - E - k$ would suffice to make it in each firm's interest to adopt the efficient technology. By assumption, since **(1)** is satisfied, society is better off as a result.

[4]The side payment can obviously take many forms, including wage or price adjustments. Of course, if transaction costs are sufficiently low to do this, this throws into question why regulation is needed at all.

[5]Another method of internalization would of course be to make industry pay for the harm it causes (either by taxation or through strict damage liability). Still another method of internalization is to create a set of marketable permits.

[6]To be sure, there may be ways for the agency to gather information about what technologies (other than those in actual use) are available to the industry, and what their costs are. Perhaps the agency can compel the industry to generate and disclose this information, or can collect it from third parties (such as firms that manufacture pollution control equipment).

III. DESIGN OF TECHNOLOGY-BASED STANDARDS

We will, therefore, confine our attention to regulatory standards that are tied to observed technologies, and are unconnected with internalization measures. We consider three "ideal types" of standards that have these features.

Firm-Specific Standards. — A firm-specific standard requires each firm to take optimal control measures *given that firm's observed capabilities.* If individual firms in the industry do not have the same capabilities, they will not be subject to the same standard. Firms with relatively inefficient control technologies will be subject to less stringent standards than firms.

Best-Technology Standards. — To set a best-technology standard, the agency identifies the optimal level of control *for the most efficient control technology observed in use.* It then requires all firms in the industry to achieve that level of control, regardless of differences in capabilities among firms. (This may be thought of as corresponding loosely to a "Best Available Technology" standard familiar in environmental law.)

Average-Technology Standards. — To set an average-technology standard, the agency examines the range of control technologies used by firms in the industry. It then identifies the *mean* location of the firms' marginal cost curves -- in effect establishing the average firm's capabilities. It sets a standard corresponding to the optimal level of control *for a firm with average capabilities.* All firms in the industry are required to achieve that level of control. (This may be thought of as corresponding loosely to standards, in environmental law and elsewhere, that use "customary" or "conventional" control technology as the benchmark for regulation.)

Figure 2 gives an example of the operation of these three standards in the model. Let *m*

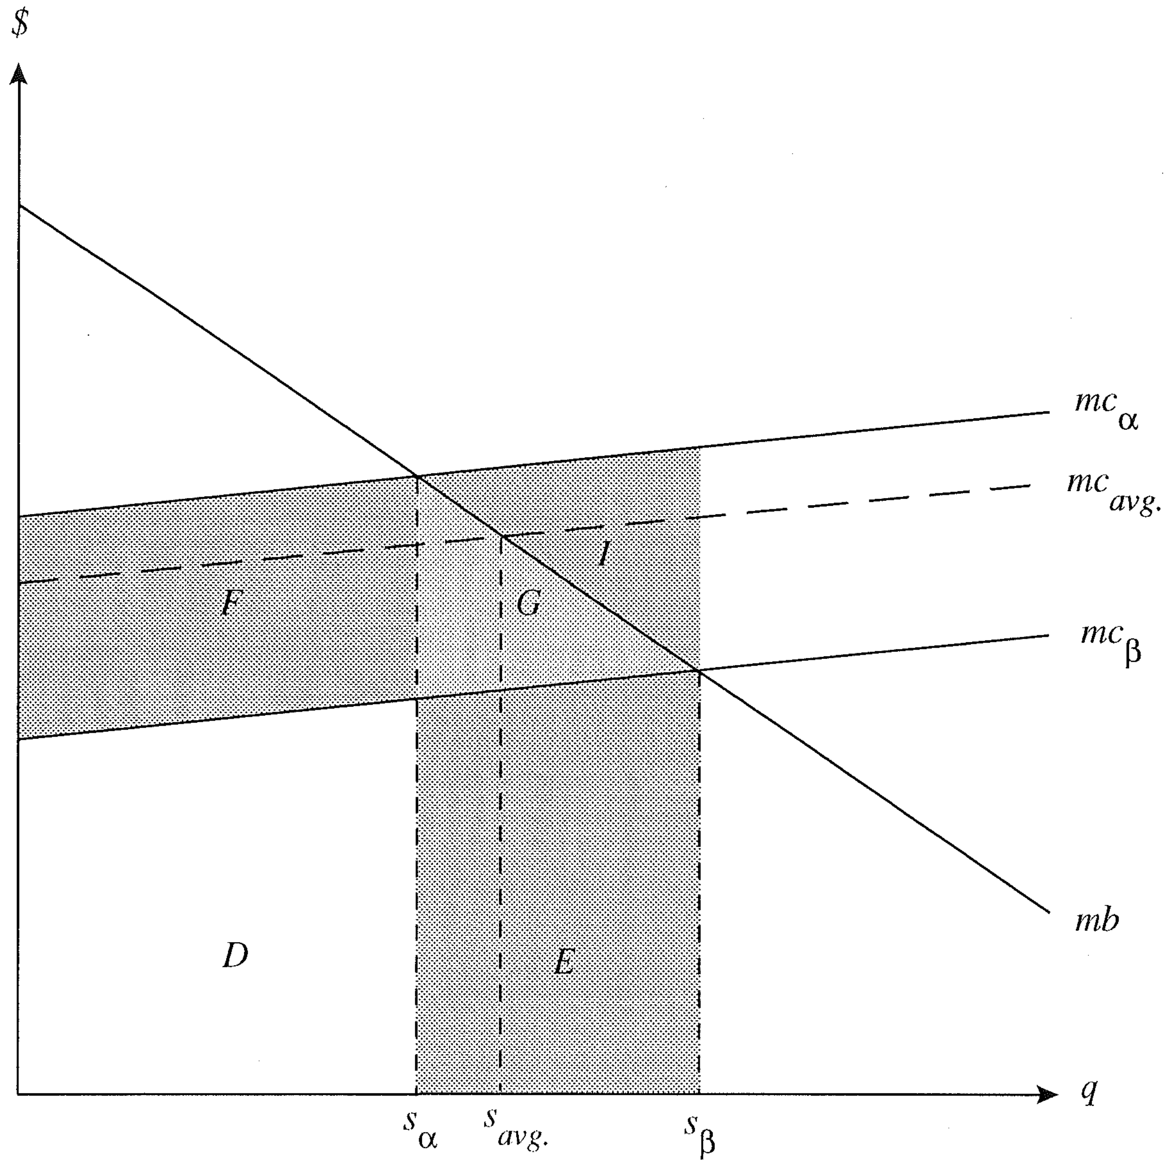

FIGURE 2. – Differences among standards when a positive fraction of firms in the industry use Beta.

represent the number of firms that use the efficient technology, and let n represent the number of firms in the industry. The line $mc_{avg.}$ depicts the *average* firm's marginal costs of control[7]; it lies the fraction m/n of the distance between the two technologies' marginal cost lines. (As I have arbitrarily drawn Figure 2, m/n is equal to about one-third.) Here is how the three standards would work under this industry configuration:

(1) Under a system of *firm-specific* regulation, the agency applies standard s_β to the firms using Beta, but applies standard s_α to the firms using Alpha. (2) Under a system of *best-technology* regulation, the agency applies standard s_β to all firms in the industry, provided that $m > 0$; if $m = 0$, then the agencies applies standard s_α to all firms. (3) Finally, under a system of *average-technology* regulation, the agency applies standard $s_{avg.}$ to all firms in the industry. Thestandard lies m/n of the distance between s_α and s_β.

IV. PAYOFFS TO INNOVATION UNDER DIFFERENT STANDARDS

Our objective is to see how firms' choice of technology is influenced by the design of regulatory standards. To do this, we treat the firms and the agency as playing a game with the following structure: each firm independently chooses a technology; the regulator then observes the firms' choice and sets a standard accordingly, depending on which the above types of standard it employs. We can think of this game as being played repeatedly; there is nothing binding about a firm's initial choice of technology. We assume, however, that the agency never *lowers* the

[7]Notice that no firm's actual costs match this average figure; everyone's true marginal cost line is either above or below the average.

TABLE 1: CONTROL COSTS UNDER DIFFERENT STANDARDS

System	Inefficient Firms' Costs	Efficient Firms' Costs
Firm-Specific	$D+F$	$k+D+E$
Best-Technology		
$m = 0$	$D+F$	-----
$m > 0$	$D+E+F+G+I$	$k+D+E$
Average-Technology	$D+F+\left(\frac{m}{n}\right)(E+G+I)^{\dagger}$	$k+D+\left(\frac{m}{n}\right)E^{\dagger}$

†Approximation.

regulatory standard in response to firms' choice of technology; standards either rise or stay the same, but never drop.

Table 1 indicates the control costs in Figure 2 sustained by each firm under the different standards, given that m firms have chosen the efficient technology.[8] This Table, however, gives only a partial representation of the payoffs from the game: since efficient firms have different costs than inefficient firms, we will need some way to express the payoff to a firm *from having relatively lower (or higher) costs than its rivals*. Let us use the following notation:

$\pi(m)$ = A firm's payoff (if it adopts Beta) attributable to having different costs than firms that use Alpha.

[8]I assume that all firms comply with the operative standard. The cost configuration would be slightly different if we assumed, for example, that firms could pay some penalty (damages or a fine) in lieu of complying. The qualitative results of the analysis would not change.

Since π varies under different standards, we use the subscripts *fs*, *at*, and *bt* to indicate firm-specific, average-technology, and best-technology regulation, respectively. Thus, for example, $\pi_{fs}(0)$ represents the value of π under best-technology standards if a firm adopts the efficient technology *when all the other firms have not (yet) adopted the efficient technology.*

We can interpret a positive value of π as meaning the firm can sell at supracompetitive prices, or being able to extract market share from inefficient rivals. I should emphasize that π should be thought of as taking into account the future behavior of other firms; in other words, it should be thought of in expected value terms, *given that other firms will behave act to maximize their own payoffs in the game*. We allow for the possibility that other firms may respond to a firm's adoption of Beta by adopting Beta themselves. If π is positive, it may represent only a *temporary* competitive advantage; the magnitude of π may depend on the ease or speed with which other firms can later switch from Alpha to Beta. Thus, for example, $\pi_{fs}(0)$ may be small if, as soon as one firm adopts the efficient technology, all the others immediately follow suit.

Our procedure will be to identify the conditions in which a given value of *m* can hold in equilibrium. Using Table 1, we examine the incentive of a single firm -- which we will call the marginal firm -- to adopt the efficient technology, given that *m* other firms have adopted the efficient technology. For *m* to hold in equilibrium, it must be that no firm can improve its position by moving from Alpha to Beta.

Firm-Specific Standards. — A given value of *m* can hold in equilibrium if and only if

$$F+\pi_{fs} < k+E. \tag{3}$$

The left-hand term represents the benefit of switching to the efficient technology, namely, the avoided costs *F* and the profit (or loss) attributable to having different costs than inefficient firms.

The right-hand term represents the costs avoided by sticking to the inefficient technology. If **(3)** obtains, the marginal firm has no incentive to deviate from the equilibrium.

Best-Technology Standards. — A value of m=0 can hold in equilibrium if and only if

$$F+\pi_{bt}(0) < k+E, \tag{4a}$$

while a value of m>0 can hold in equilibrium if and only if

$$F+G+I+\pi_{bt} < k. \tag{4b}$$

The left- and right-hand terms here have the same interpretation as in **(3)**.

Average-Technology Standards. — A given value of m can hold in equilibrium if and only if

$$F + \left(\tfrac{m}{n}\right)(G+I) +\pi_{at} < k+\left(\tfrac{1}{n}\right)E. \tag{5}$$

The left- and right-hand terms here have the same interpretation as in **(3)**. The term $\left(\frac{1}{n}\right)E$ reflects the fact that the standard only moves up a bit in response to the marginal firm's adoption of Beta.

V. TECHNOLOGY CHOICE IN EQUILIBRIUM

Using the above expressions, we are now in a position to examine the conditions under which the industry will adopt the efficient technology.

A. *Firm-Specific vs. Uniform Standards*

Let us begin by comparing firm-specific standards to the two types of uniform standard.

Under *firm-specific* standards, the entire industry will retain the inefficient technology unless

$$F > k+E. \tag{6}$$

In contrast, under *best-technology* standards, the entire industry will adopt the efficient technology if

$$F > k+E-\pi_{bt}(0). \tag{7}$$

And under *average-technology* standards, the entire industry will adopt the efficient technology if

$$F > k+\left(\frac{1}{n}\right)E. \tag{8}$$

From these expressions we can derive the following claim.

> PROPOSITION 1. Uniform standards, but not firm-specific standards, may induce industry to adopt the efficient technology even if the effect is to raise industry's control costs.

The proof is left to the Appendix. To see the point informally, observe that **(6)** is the same as **(2)**, meaning that firm-specific standards will induce adoption of the efficient technology if and only if the effect is to lower industry's costs. Yet **(7)** and **(8)** may be satisfied even when **(6)** is not, meaning that uniform standards will sometimes induce adoption of the efficient technology even when its adoption is not in industry's collective interest.

Why do uniform standards have this effect? Consider best-technology standards. Their essential feature is they *tend to impose a heavy penalty on relatively inefficient firms*, creating a corresponding competitive advantage for relatively inefficient firms. If **(7)** holds, then each firm will seek to get a competitive advantage -- albeit a temporary one -- over its rivals, even if its own costs rise as a result. (Notice that no firm actually *succeeds* in gaining at the other's expense; in this model, all firms adopt the efficient technology. But this does not stop them from trying; any firm that unilaterally stuck to the inefficient technology would be shooting itself in the foot.)

Turn now to average-technology standards. These also tend to put relatively inefficient firms at a competitive disadvantage; but in addition, they *tend to encourage adoption of the efficient technology by assuring that no single firm's action will do much to affect the regulatory standard.* If n is sufficiently large, each firm knows that *its own* adoption of the efficient technology will scarcely affect the standard; this makes it more likely that adopting the efficient technology will lower its overall control costs. If **(8)** holds, each firm, wanting to lower its own costs, adopts the efficient technology, even though the net effect may be to raise everyone's costs.

B. *Average- vs. Best-Technology Standards*

Which type of uniform standard furnishes the greater stimulus to adopt the efficient technology? To examine this question, we will compare the circumstances in which $m=0$ can hold in equilibrium. If $m=0$ is an equilibrium under one standard but not under the other, we will say that the latter standard creates the greater stimulus to innovation.[9] We have the following result:

> PROPOSITION 2. Depending on the circumstances, either type of uniform standard may furnish a greater stimulus for industry to adopt the efficient technology. In particular:
>
> (*a*) When firms have difficulty maintaining a technological edge over rivals, then average-technology standards are likely to furnish the greater stimulus to innovation.
>
> (*b*) When firms can maintain a technological edge over rivals, then neither

[9]More precisely, if $m=0$ is an equilibrium under one standard anytime it is an equilibrium under the other standard, but not vice versa, then we will say the latter standard creates the greater stimulus to innovation.

type of uniform standard clearly furnishes the greater stimulus to innovation.

The proof is left to the appendix. The central point can be captured by the following inequality:

$$\left(\frac{n-1}{n}\right)E \quad > \quad \pi_{bt}(0) - \pi_{at}(0)\,. \tag{9}$$

If **(9)** holds, then an average-technology standard furnishes the greater stimulus to innovation; if **(9)** does not hold, then a best-technology standard creates the greater stimulus.[10]

The interpretation of **(9)** is as follows. Assume a firm introduces the efficient technology when when m=0. On the one hand, from the firm's standpoint, the average-technology standard generates *lower compliance costs* than the best-technology standard; the left-hand side of **(9)** captures the difference between the two standards in this regard. On the other hand, the best-technology standard generates a *larger cost differential* between efficient and inefficient firms -- and thus confers a *greater competitive advantage* on efficient firms -- than does the average-technology standards; this effect is captured by the right-hand side of **(9)**.[11] The question is which effect predominates.

Suppose that firms can easily switch to the efficient technology, so that efficient firms have difficulty maintaining a cost advantage over inefficient rivals. Then the right-hand side of **(9)** is

---

[10]More precisely: If **(9)** holds, then m=0 will be an equilibrium under best-technology standards anytime it is an equilibrium under average-technology standards, but not vice versa.

Conversely, if **(9)** does not hold, then m=0 will be an equilibrium under average-technology standards anytime it is an equilibrium under best-technology standards, but not vice versa.

[11]We assume for purposes of this discussion that efficient firms have lower overall costs than inefficient firms, so that both $\pi_{at}(0)$ and $\pi_{bt}(0)$ are nonnegative.

likely to be small. For under best-technology standards, once a single firm introduces the efficient technology, all other firms in the industry have an incentive to follow suit as soon as possible; if they can quickly do so, they will. The competitive advantage enjoyed by the firm to introduce Beta will thus be short-lived. As the ease of switching technologies goes up, the value of $\pi_{bt}(0)$ approaches zero, meaning that **(9)** is likely to be satisfied.

Suppose, however, that switching technologies is difficult or time-consuming, so that efficient firms can maintain a cost advantage over inefficient rivals for a substantial period of time. Here it is difficult to say much about the right-hand side of **(9)**. For as switching technologies becomes more difficult, both $\pi_{bt}(0)$ and $\pi_{at}(0)$are likely to grow large. Nonetheless, in some instances the *difference* between them may stay sufficiently small to satisfy **(9)**; in other instances, the difference between them may grow so large that **(9)** is not satisfied. One cannot a priori rule out either possibility.

C. *Private vs. Social Incentives*

Let us now consider whether firms' incentives are aligned with society's under the different types of standard. Firms' incentives to adopt a more efficient technology are *insufficient* if they would refrain from adopting it, even though its efficiency gains (when the technology is optimally used) exceed its costs of adoption. Conversely, firms' incentives are *excessive* when they wouldadopt a new technology, even though its efficiency gains (when the technology is optimally used) are less than its costs of adoption. We have the following result:

> PROPOSITION 3. Under uniform standards, firms' incentives may be either excessive or insufficient. Under firm-specific standards, firms' incentives may be insufficient, but are

never excessive.

The proof is left to the Appendix; the intuition is as follows. Under firm-specific standards, a firm's only inducement to innovate is the prospect of lowering its own control costs. As a result, industry (as we have seen) will only adopt the efficient technology if -- optimally used -- it lowers industry's overall control costs. But anytime this condition obtains, society is better off from the efficient technology's adoption, since control costs constitute social costs.

In contrast, under uniform standards, a firm potentially has a second inducement to innovate: to get a competitive advantage over its rivals by *lowering its costs relative to theirs.* This second motivation may lead a firm to innovate even when its own costs go up: so long as its rivals' costs go up even more than its own, the firm may be better off as a result. This jockeying for competitive advantage may lead to innovations whose costs exceed their social benefits.

Under average-technology standards, the crucial test is whether

$$\pi_{at}(0) > G + \left(\tfrac{1}{n}\right)E. \tag{10}$$

If this inequality holds, some firms may adopt the efficient technology even though society would prefer none did. If the inequality does not hold, then we have the opposite problem: some firms may refrain from adopting the efficient technology, even though society would prefer all firms adopted it. Under best-technology standards, the analogous test is whether

$$\pi_{bt}(0) > G + E. \tag{11}$$

If this holds, firms' incentives may be excessive; if it does not hold, firms' incentives may be insufficient.

As expressions **(10)** and **(11)** imply, it is *the payoff from having a competitive advantage over rivals* that may give firms excessive incentives to adopt new control technologies. Observe that if the π terms are small -- meaning either that costs advantages are hard to maintain or that they are not of much value -- then these expressions are unlikely to be satisfied, meaning firms' incentives will not be excessive. If the π terms are sufficiently large, however, then industry will rationally invest in new technology even when the effect is to raise not only its own overall costs but society's as well.

V. APPLICATIONS AND EXTENSIONS

The ideal-type standards we have examined capture two basic choice problems in designing regulations that are tied to firms' observed capabilities.

Uniformity vs. Differentiation. — To what extent will standards will distinguish between firms on the basis of their differing control capabilities? At one extreme, the regulator could in principle set a single uniform standard for the industry, which all firms would have to comply with regardless of their individual control costs or other unique characteristics. At the other extreme, the regulator could set standards on a firm-by-firm basis, determining what level of emissions control is appropriate for each firm and setting a standard accordingly. As an intermediate approach, the regulatory could divide the industry into groups of firms, and set a standard for each group.

In practice, the extent of regulatory differentiation among firms is likely to fall between these two extremes. Setting an individual standard for each firm, based on a detailed examination of that capabilities, would involve prohibitive administrative costs.[12] On the other hand, setting and enforcing a single uniform standard for an entire industry is frequently also infeasible.[13] Yet there is, in principle, considerable maneuvering room between these extremes. There can be anywhere from a lot of differentiation among firms (creating many categories of firms, giving each its own standard) to only a little.[14]

The model gives some insight into the consequences of moving in one or the other direction. The more an agency differentiates among firms -- the more categories it creates, the

[12]Strictly speaking, there will always be an element of generalization in even the most individualized possible regulatory system. (For example, generalizations of the sort, "machine *X* is capable of such-and-such performance"). Eliminating all generalizations would be infinitely costly. But generalization implies some uniformity of treatment (for example, a refusal to inquire into potential variations among units of machine *X*).

[13]For example, political pressures may force the agency to give a "break" to less capable or less profitable firms, lest such firms be shut down (with ensuing job losses and the like) as a result of regulation.

[14]Differentiation can take the form of either initially establishing separate categories, or instead beginning with a single category but allowing firms to seek (formal or de facto) variances. Also, standards that are nominally uniform may turn out not to be so in enforcement. (Suppose the agency concentrates most of its enforcement resources on the violators who could most easily have complied with the standard, where the effect of compliance is least likely to result in job losses, and so forth.) There is some anecdotal evidence of this in environmental law.

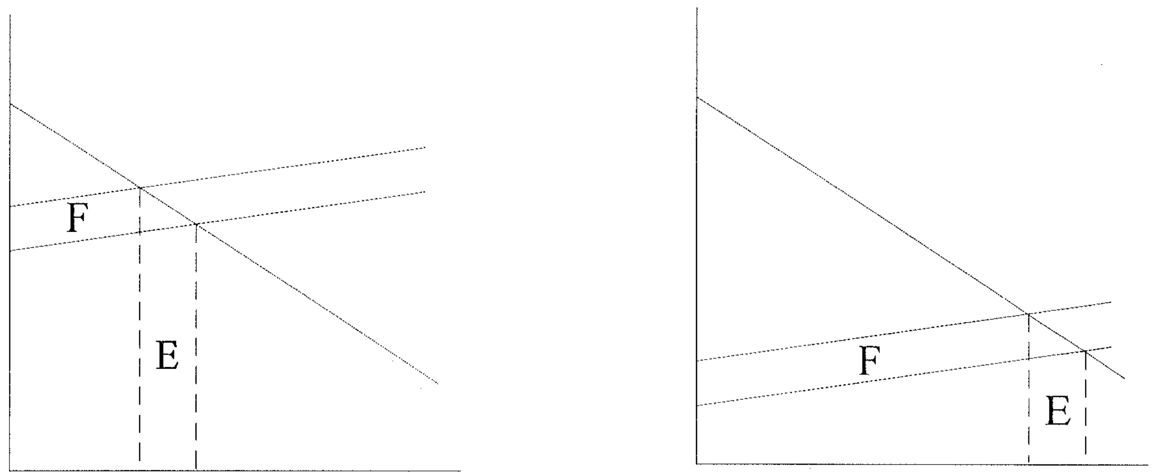

FIGURE 3. — Lowering the marginal cost curves changes the relative sizes of E and F, increasing the likelihood that expression **(2)** is satisfied.

smaller n becomes in each category -- the less likely innovation becomes. Under highly differentiated standards, innovation is then unlikely unless the new technology, optimally used, gives the firms in each category lower overall costs than the old technology. The model also gives us some rough insight into when this condition obtains: generally speaking, the greater the marginal cost of control, the less attractive innovation will be under highly differentiated standards. (See Figure 3.)

Content of Uniform Standards. — To the extent it employs uniform or categorical standards, the regulator must decide on a baseline for regulating firms with different capabilities. If firms have different costs of control, the optimal degree of emissions control varies from firm to firm. But by assumption, the agency will only set a single standard for the entire category. What should the standard require?

Regulators confront this problem all the time in environmental law. Consider Congress' choice between "best available technology" (BAT) and "best conventional technology" (BCT) requirements in pollution statutes. One interpretation of these formulations is the reference point the agency is instructed to use in setting standards; under BAT, the agency is to instructed to look at what is done — or what level of performance is achieved — by the firms in the upper (more effective) end of the category; under BCT, the agency is instructed to look at the firms more toward the middle. Analogous reference point questions come up in other regulatory areas.[15]

There is obviously no sharp dividing line between what I have called best-technology and average-technology standards; they represent extremes on a continuum. In constructing an average-technology standard, for example, the agency can use as its reference point the mean of all firms in the category, or instead the mean of only (some fraction of) the best-performing firms in the category; in the latter case an average-technology standard converges on a best-technology standard. Again, however, there is maneuvering room between the two extremes, and the model tells us the impact of moving in either direction. The more the standard is geared to the top-performing fraction of firms in a given category, the greater the power of an individual firm to raise (by adopting a more efficient technology than its rivals) the standard imposed on all

---

[15]For example, OSHA must choose between standards based on "national consensus" (meaning industry custom) and standards based on more stringent criteria. In tort law, courts confront analogous problems in (1) defining the standard of care under negligence rules (what is the role of custom?) an (2) giving content to notions such as "defectiveness" and "state of the art' under doctrines of strict liability for product-related injuries.

firms in the category.[16] If cost advantages are easy to maintain -- if rivals have a hard time quickly matching the efficient firm's choice of technology -- then moving in the direction of a best-technology standard may encourage greater innovation. (Expression (9) may become harder to satisfy.)

If, instead, cost advantages are difficult to maintain -- perhaps because rivals can easily match the efficient firm's choice of technology, or (equivalently) because entry barriers are low -- then moving in the direction of an average-technology standard will probably encourage greater innovation. (Expression (9) becomes easier to satisfy.) For when relative cost advantages over rivals are hard to maintain, firms will care mostly about minimizing their control costs in absolute terms. As we have seen, average-technology standards, being comparatively insensitive to individual firms' choice of technology, make it most likely that adopting a more efficient technology will lower firms' control costs.

From the standpoint of social welfare, the model yields no general prescription for the optimal design of technology-based standards. As we have seen, applying different standards to different firms risks producing insufficient incentives for innovation; using (more) uniform standards reduces that risk, but creates the opposite risk of excessive incentives. Perhaps the only clear prescription is that the regulator who is constrained to employ technology-based standards must be attentive to both risks.

---

[16] If the agency looks at the average costs of all firms in a large category, a given firm's choice of technology only affects the relevant average a little bit. If, instead, the agency looks at the average costs of (say) the top three firms in the category, then theaction of one of those three firms can affect the relevant average a lot. The general point in the text assumes the firms' costs follow a roughly normal distribution.

APPENDIX

Proof of Proposition 1

Begin with firm-specific standards. For a given firm, sticking with Alpha yields control costs $D+E$, while switching to Beta yields control costs $k+D+E$; these payoffs are independent of what any other firm does at any time. Thus, regardless of what any other firm's actual or anticipated actions are, a firm will not adopt Beta unless

$$F > k + E. \qquad \textbf{(A1)}$$

Accordingly, if **(A1)** is not satisfied, no firm in the industry will adopt Alpha.

Now consider best-technology standards. Suppose that

$$F + \pi_{bt}(0) > k + E, \qquad \textbf{(A2)}$$

meaning that at least one firm will adopt Beta. Then it must also be true that

$$F+G+I + \pi_{bt}(m) \;>\; k + E, \qquad \textbf{(A3)}$$

for $m>0$ -- meaning that if one firm adopts Beta, all others will do so as well. To see this, observe that $\pi_{bt}(m)$ is nonnegative (for any m) iff efficient firms have lower total control costs than inefficient firms, which (Table 1 shows) is true when

$$F+G+I \;>\; k; \qquad \textbf{(A4)}$$

thus, **(A2)** implies **(A4)**, which implies **(A3)**. Thus, **(A2)** suffices to ensure that the entire industry adopts Beta. Finally, observe that **(A2)** may hold while **(A1)** does not -- meaning that the entire industry may adopt Beta even though doing so is not in its collective interest.[17]

Now consider average-technology standards. Suppose that

---

[17]Recall that **(A1)** is the same as **(2)** in the text, which furnishes the test for whether adoption of the new technology is in industry's interest.

$$F > k+\left(\tfrac{1}{n}\right)E. \tag{A5}$$

Then it follows that for any value of m,

$$F + \left(\tfrac{m}{n}\right)(G+I) + \pi_{at} > k+\left(\tfrac{1}{n}\right)E. \tag{A6}$$

To see this, observe that (from Table 1) $\pi_{at}(m)$ is positive iff

$$F+\left(\tfrac{m}{n}\right)[G+I] > k; \tag{A7}$$

so **(A5)** implies **(A7)**, which implies **(A6)**. But if **(A6)** holds, no value of $m < n$ can hold in equilibrium. Thus, **(A5)** suffices to ensure that the entire industry adopts Beta. Finally, observe that **(A5)** may hold while **(A1)** does not.

Proof of Proposition 2

Under a best-technology standard, $m=0$ can hold in equilibrium unless

$$F+\pi_{bt}(0) > k+E; \tag{A8}$$

while under an average-technology standard, $m=0$ can hold in equilibrium unless

$$F + \pi_{at}(0) > k+\left(\tfrac{1}{n}\right)E. \tag{A9}$$

Now, expressions **(A2)** and **(A3)** establish that under best-technology standards, if a single firm adopts Beta, all the rest will follow. Suppose, therefore, that firms are capable of switching technologies instantaneously. Then a firm that introduces Beta when $m=0$ will not obtain a cost advantage over inefficient rivals, because the latter will immediately switch to Beta themselves.

Thus, $\pi_{bt}(0)$will approach zero. But if $\pi_{bt}(0)$ is zero, then satisfaction of **(A8)** implies satisfaction of **(A9)**. (To see this, observe that $F > k+E$ implies $F > k$, which (by **(A7)**) implies that $\pi_{at}(0) \geq 0$, meaning that **(A9)** must hold.) Thus, if $\pi_{bt}(0)$ is zero, then $m=0$ is an equilibrium under a best-technology standard anytime it is an equilibrium under an average-technology standard, but not vice versa.

Proof of Proposition 3

Expression **(1)** is the test for whether society is better off from the introduction of the efficient technology. If **(1)** holds, society's preferred outcome is to have all firms adopt the efficient technology; if **(1)** does not hold, society's preferred outcome is to have no firms adopt it.[18] Thus, we want to see how well firms' incentives to adopt the efficient technology map onto expression **(1)**.

Consider firm-specific standards. By inspection, **(6)** implies **(1)**, but not vice versa. Thus, any time industry adopts the efficient technology, society is better off as a result; however, industry will sometime refrain from adopting the efficient technology even though society would prefer its adoption.

Now consider average-technology standards. If

$$\pi_{at}(0) > G+\left(\tfrac{1}{n}\right)E, \tag{A10}$$

[18]The latter point does not depend on the assumption that firms in fact make optimal use of the efficient technology. Expression **(1)** determines whether the new technology, optimally used, makes society better off. If the answer is no, then by definition the new technology cannot make society better off if it is *not* optimally used. Thus, if **(1)** does not hold, society does not want the new technology introduced.

then it is is possible to choose a value of k for which neither **(1)** nor **(5)** holds (at m=0); thus, at least some firms will adopt the efficient technology even though society is better off it no firms adopt it. If the above expression does not hold, then it is possible to choose a value of k for which both **(1)** and **(5)** hold, so that some firms may refrain from adopting a technology that society would prefer the entire industry adopt.

Finally, consider best-technology standards. If

$$\pi_{bt}(0) > G+E, \qquad \textbf{(A11)}$$

then it is possible to choose a value of k for which neither **(1)** nor **(4a)** holds, so that industry will adopt the efficient technology even though society would prefer that no firm adopt it. Conversely, if the above expression is not satisfied, it is possibile to choose a value of k for which both **(1)** and **(5a)** hold, meaning that industry may refrain from adopting a technology society would prefer that it adopt.