

EVALUATING INSTRUMENTS OF ENVIRONMENTAL POLICY: A COMMENT ON PROFESSOR RICHARDS

NATHANIEL O. KEOHANE*

I. INTRODUCTION

In his article, "Framing Environmental Policy Instrument Choice,"¹ Professor Richards proposes a taxonomy of environmental policy instruments as a guide in drawing comparisons between the instruments. His logical apparatus provides a useful framework for organizing environmental policies and offers a measure of insight about the relationships among them. Professor Richards is less successful, however, in evaluating the policy instruments, largely because of his attempt to apply transaction cost economics to environmental policy.

In this Comment, I discuss four points in turn. First, Professor Richards frames his discussion in terms of a stylized cost-minimization problem. Although a useful first step, his presentation of the problem as a *static* one limits his evaluation of the policy instruments with respect to the dynamic effects on technological change. Framing the problem in explicitly dynamic terms suggests several ways in which the cost components and even the constraints of the problem change over time.

Second, Professor Richards draws a distinction between policy instruments that require payment from polluters to the government and those that do not. He argues that whether or not private firms pay for their emissions determines who bears the burden of environmental damages, and he returns to this distinction in evaluating the instruments on the grounds of "equity and environmental justice." This discussion, however, confuses the incentive properties of emissions payments with equity considerations. Such payments do not compensate the victims of pollution; indeed, they are not necessarily even *equal* to the damages from pollution.

* Ph.D. Candidate, Political Economy and Government, Harvard University.

1. Kenneth R. Richards, *Framing Environmental Policy Instrument Choice*, 10 DUKE ENVTL. L. & POL'Y F. 221 (2000).

Third, Professor Richards overlooks an important class of policy instruments: ambient taxes and emissions permits. Indeed, the distinction Professor Richards makes between controlling inputs and controlling outputs can be extended to the choice of the “locus of regulation”—that is, at what stage to regulate, along the chain from inputs to emissions to eventual environmental damages.

Finally, Professor Richards devotes considerable attention to what he calls the “governance costs” associated with implementing environmental policy instruments. Here he draws heavily on transaction cost economics, particularly the work of Oliver Williamson. Arguing that the conditions of uncertainty and asset specificity apply equally well to the environmental policy arena as to contractual relationships between private parties, Professor Richards concludes that, in some cases, market-based policy instruments may invite government opportunism, raising their true social costs and leading to a preference for command-and-control regulation.

However, it is by no means clear that the environmental policy arena is analogous to the contractual setting envisioned by the transaction cost theorists. The examples provided by Professor Richards are better illustrations of a cost-revelation problem than the “hold-up” problem that is the focus of the transaction cost literature. I suggest a hypothetical hold-up problem in environmental policy and argue that its relevance rests on special assumptions about the regulator’s treatment of firms, which do not generally apply to market-based instruments. Moreover, even to the extent that “governance costs” arise in environmental regulation, hierarchical control is unlikely to represent a solution.

II. THE COST-MINIMIZATION PROBLEM

Economic theory is much more comfortable describing efficiency than equity: the former criterion mobilizes the powerful apparatus of cost-benefit analysis, while the latter involves the messy complications of interpersonal comparisons. In evaluating environmental policy instruments, therefore, economists have tended to give greatest weight to efficiency concerns—or, in the absence of information on benefits, to cost-effectiveness concerns. Professor Richards follows suit. He proposes the following cost-minimization problem as a starting point for evaluating policy instruments:

Minimize (PC + IC + TX)

Subject to the following constraints:

Pollution abatement requirement

Legal constraints

Political constraints.²

In his formulation, PC denotes “production costs,” or the costs of pollution abatement; IC denotes “implementation costs,” notably monitoring costs and what Professor Richards terms “governance costs”; and TX denotes public finance impacts, or the deadweight loss from environmental regulation due to pre-existing distortionary taxes.³

As stated, however, this problem lacks an explicit time dimension. To be sure, at certain points in the discussion Professor Richards recognizes the dynamic aspects of the policy problem.⁴ But, he unduly narrows his focus by framing the problem in static terms.

With regard to pollution abatement costs, two dynamic issues stand out. First, different policy instruments allow different degrees of discretion in the timing of pollution abatement. Indeed, Professor Richards includes this as the sixth “dimension” of his taxonomy.⁵ Allowing banking of tradeable emissions permits, for example, gives polluting firms the flexibility to allocate pollution control efforts across time as well as space. In theory, this allows intertemporal cost-minimization, as firms can equalize marginal abatement costs at different time periods—for example, by banking permits when costs are low for use later when marginal abatement costs increase (perhaps due to a reduction in the number of new permits granted by the regulator). A shift in emissions from earlier to later periods also enhances environmental benefits, since future damages are discounted relative to present ones. Both effects of banking, for example, are apparent in the experience thus far with the tradeable allowances provision of Title IV of the 1990 Clean Air Act Amendments.⁶

A second, and perhaps more important, dynamic element in environmental policy concerns the effects of policy instrument choice on the innovation and diffusion of new pollution control technologies.

2. See Richards, *supra* note 1, at 228.

3. See *id.*

4. See *id.* at 228, 248, 255.

5. See *id.* at 248.

6. See A. DENNY ELLERMAN, ET AL., MARKETS FOR CLEAN AIR: THE U.S. ACID RAIN PROGRAM (forthcoming 2000)(manuscript at 58, on file with author).

Dynamic incentives are presumably especially important for recently-appearing environmental policy issues, such as global climate change, that act over long time scales. A growing body of literature in environmental economics (part of which Professor Richards cites) concludes generally that market-based instruments provide greater incentives for innovation and diffusion of policy instruments than do more rigid approaches, such as performance or technology standards, although the question of *efficient* incentives for technological change remains largely an open one, dependent on details such as the form of the pollution damages function.⁷ The dynamic effects of policy instrument choice on technological change are also implicit in Professor Richards' concerns about regulatory opportunism.⁸

The importance of a dynamic framework also extends to the other two elements of Professor Richards' constrained minimization problem: implementation costs and public finance effects. Monitoring costs are presumably highest at the beginning of a regulatory effort: as both the regulator and regulated firms learn about the new policy and about the technologies used for monitoring and measurement, those costs are likely to fall. Moreover, the bulk of monitoring costs may involve capital investments in new devices, such as the Continuous Emissions Monitoring Systems (CEMS) used to measure SO₂ emissions under the tradeable allowances program. Thus, the implementation costs for a particular instrument are likely to vary over time. Similarly, "governance costs" may be time-dependent to the extent that they depend on uncertainty about proper regulatory stringency, the regulator's preferences, or private abatement costs.⁹

Finally, one may even view the *constraints* of Professor Richards' characteristic problem as essentially dynamic. The pollution abatement requirement, of course, may vary over time—not only due to changes (perhaps unforeseen) in preferences or costs, but also as part of a deliberate and pre-announced gradual reduction of pollution levels or expansion of coverage. For example, the SO₂ tradeable allowances program has been implemented in two phases: Phase I, beginning in 1995, included only the dirtiest power plants (the largest

7. See, e.g., Paul B. Downing & Lawrence J. White, *Innovation in Pollution Control*, 13 J. ENVTL. ECON. & MGMT. 18 (1986); Scott R. Milliman & Raymond Prince, *Firm Incentives to Promote Technological Change in Pollution Control*, 17 J. ENVTL. ECON. & MGMT. 247 (1989); Carolyn Fischer, et al., *Instrument Choice for Environmental Protection When Technological Change is Endogenous*, RESOURCES FOR THE FUTURE DISCUSSION PAPER 99-04 (1998).

8. These effects are discussed extensively *infra* Part V.

9. For more on the temporal element to governance costs and a critique of Richards' approach, see *infra* pp. 401, 410-411.

emitters of sulfur dioxide), while Phase II (beginning in 2000) extends the program to a much wider universe of polluters. Moreover, the number of annual allowances (which can be used any time in or after their “vintage year”) falls in Phase II, representing a “ratcheting” of pollution control.

Political and legal constraints may also change—and, importantly, that change may be endogenous to the choice of instrument. Casual observation suggests that the past two decades have seen a relaxation of opposition to market-based instruments. At least some of this change of heart might reasonably be ascribed to gradual but growing experience with such instruments—starting with the emissions trading provisions of earlier incarnations of the Clean Air Act, continuing with the lead phase-down program of the mid-1980s, and culminating in the success of the SO₂ allowances program.¹⁰ Indeed, in some sense, political constraints can be seen as a function of the “stock” of experience with a particular instrument, or class of instruments. A similar argument could apply to legal constraints.

In place of Professor Richards’ static problem, therefore, we might better view the representative constrained optimization problem (minimizing environmental costs, E) in explicitly dynamic terms:

$$\text{Minimize } \{x_t\} E \left[\int_0^{\infty} (PC_t(x_t) + IC_t(x_t) + TX_t(x_t)) e^{-\delta t} dt \right]$$

subject to : $Q(t) \leq \bar{Q}(t)$ [abatement requirement]

$x_t \in X(t)$

$X(t) = F(\text{politics, legal constraints, experience, etc.})$

This formulation expressly recognizes the dynamic aspects discussed above, including the possible evolution over time of the “politically and legally feasible set,” here denoted by $X(t)$ and assumed to be a function of politics, legal constraints, experience with various policy instruments, and so on. Uncertainty is prominent: the regulator seeks to minimize the *expected* costs, given uncertainty. And the (social) discount rate δ —itself a product of social deliberation—is presented explicitly as a parameter of the policy problem.

10. For a related discussion of the importance of the learning curve for the acceptance of market-based instruments (especially as it relates to legislators), see Nathaniel O. Keohane, et al., *The Choice of Regulatory Instruments in Environmental Policy*, 22 HARV. ENVTL. L. REV. 313, 357-63 (1998).

Professor Richards does not ignore the dynamic aspects of policy instrument choice in his discussion. But he integrates these aspects only in an *ad hoc* manner. Incorporating dynamics directly into the cost-minimization problem reveals the range of ways in which both the costs and the constraints of environmental regulation vary over time, in part as a function of accumulated experience and past regulation.

III. POLLUTER PAYMENTS AND ENVIRONMENTAL DAMAGES

The organizing theme of Professor Richards' taxonomy is the degree of government involvement in pollution control, along a range of relevant dimensions. One such dimension is the "distribution of abatement and environmental damage costs."¹¹ Professor Richards argues that policy instruments differ not only by the extent to which private polluters pay the costs of abatement, but also by the extent to which those polluters bear the damages from pollution. He has in mind policies, such as auctioned tradeable permits or emissions taxes, that require polluters to pay for their emissions. In his conclusion, he returns to this issue, arguing that "the distinction between abatement costs and residual environmental damages emphasizes the issues of equity and fairness inherent in deciding who bears these two types of costs."¹²

The connection between payments by polluters and distributional equity is tenuous for two reasons. First, emissions payments need not equal environmental damages. Second, and more fundamentally, payments by polluters do not represent compensation to the victims of pollution. To see the first point, consider a zero-baseline emissions tax (the same principles apply directly to a system of auctioned tradeable permits). Suppose that the costs of pollution are described by an upward-sloping marginal damages function, and suppose (perhaps due to political constraints, or imperfect knowledge of the marginal damages function) that the regulator is restricted to using a linear tax—one for which the tax on any given unit of pollution is constant. Then the tax payment from the regulated industry will *exceed* the damages from pollution. This simple result is illustrated in Figure 1, which plots marginal pollution damages against marginal (industry) abatement costs. Given a tax on emissions equal to P , polluters (in the aggregate) emit Q units of emissions (the level at

11. Richards, *supra* note 1, at 237.

12. *See id.* at 282.




Figure 1. Abatement costs and environmental damages.

which aggregate marginal abatement costs equal the tax). Total environmental damages are given by the area of the triangle OAQ , and total tax revenue by rectangle $OPAQ$. Of course, in the special case where marginal damages are constant, the tax payments and damages from pollution coincide, but such agreement occurs only by the accident of the damages function.¹³

Whether or not the tax payment happens to equal pollution damages, however, the victims of pollution do not necessarily receive compensation. Indeed, economic theory suggests that Pareto optimality requires that victims *not* be compensated—at least, not in proportion to the damages they suffer. Any compensation paid to the victims must be lump-sum, rather than a function of the damages borne.¹⁴ Rather than serving to “punish” polluters and compensate victims, emissions payments provide an incentive to reduce emissions by attaching a price to pollution on the margin. Because every polluter faces the same marginal tax rate, they independently choose abatement levels at which their marginal abatement costs are equal, satisfying a necessary condition for cost-effectiveness. That is, emissions payments encourage an efficient allocation of pollution con-

13. More generally, if marginal damages are known and the regulator can levy a nonlinear tax, then a tax schedule equal to the marginal damages function will result in tax payments equal to damages. But environmental economists typically assume that the complexity of a nonlinear tax—in particular, the need to know the marginal damages schedule—rules out its use by the regulator. Indeed, restricting the regulator to a linear tax (i.e., a constant tax per unit of emissions) seems particularly fitting in the cost-effectiveness framework provided by Professor Richards. The emphasis on cost-effectiveness rather than efficiency presumably reflects an implicit assumption that the regulator lacks knowledge of the true damages from pollution.

14. See WILLIAM J. BAUMOL & WALLACE E. OATES, *THE THEORY OF ENVIRONMENTAL POLICY* 42-44 (2d ed., 1988).

trol.¹⁵ Distributional equity, on the other hand, would be served by a distinct program of compensatory payments, which—like other redistributive measures—imposes a cost in efficiency terms.¹⁶ From the point of view of economic theory, making the polluter pay has little to do with equity and everything to do with efficiency.

Moreover, Professor Richards assumes that taxes and auctioned permits are revenue-neutral, so that revenues are “recycled” to reduce other distortionary taxes.¹⁷ Revenue neutrality is inconsistent with victim compensation: proceeds from an environmental tax cannot simultaneously reduce distortionary taxes and compensate for pollution damages. Compensating the victims, however, would seem to be central to any notion of “fairness” or “justice” involved in making polluters pay. Merely raising the costs to polluters, without attending to those who suffer from the pollution, would seem a meager kind of justice.

To clarify the distinction between emissions payments and fairness, consider the following revenue-neutral emissions tax: polluters pay a charge on each unit of emissions, and then the tax revenue is returned to the polluting industry in the form of lump-sum payments or an industry-wide reduction on corporate income taxes. Such a tax scheme would be cost-effective and would require polluters to pay on inframarginal as well as marginal emissions, yet it would not represent an improvement in “equity” or fairness. Thus, the mere fact that firms pay for their emissions has no independent distributional significance.

Professor Richards equates a transfer of resources from polluters to the government as a payment for environmental damages. But, consider the implications of this argument for policy instruments, such as “grandfathered” (i.e., freely allocated) tradeable permits, that involve transfers from government to polluters. The value of such a transfer is equal, in principle, to the amount that polluters would pay to the government under an otherwise equivalent system of auctioned permits. However, one would not conclude that grandfathered permits represent a payment of environmental damages to the polluter. Such transfers between government and private industry do not affect the distribution of the environmental damages themselves.

15. If the marginal damages function is known, and the tax or number of permits set accordingly, these instruments also induce an efficient level of pollution.

16. See generally Sally Holtermann, *Alternative Tax Systems to Correct for Externalities, and the Efficiency of Paying Compensation*, 43 *ECONOMICA* 9 (1976).

17. See Richards, *supra* note 1, at 241.

In general, therefore, whether or not polluters pay for their emissions does not affect who bears the actual environmental damages—the health risks from dirty air, say, or the loss of riparian habitat in polluted streams. Regardless of the policy instrument chosen, or the size of the transfer from private industry to government, those damages are borne by society. A compensation scheme would be logically distinct and could be added equally well to any policy instrument.

Professor Richards would be correct in distinguishing among instruments on the basis of whether private industry pays for their emissions. But while this distinction is important on efficiency grounds, it is not significant on the grounds of distributional equity.

IV. AMBIENT TAXES AND PERMITS AND THE “LOCUS OF REGULATION”

While he includes most of the major policy instruments in his taxonomy, and even suggests a few new ones, Professor Richards overlooks an important class of approaches to environmental policy that are particularly relevant for pollution control: *ambient* taxes and permits, which are tied to pollution concentrations at specific receptors rather than to emissions. Under such a system, polluters face a “price on pollution concentration” resulting from emissions, rather than a price on emissions.¹⁸ While ambient permit systems are largely of theoretical interest, some scholars have proposed their use in federal environmental policy in the United States.¹⁹

Under an ambient tax or permit system, the marginal price on emissions varies among polluters according to their location. Emissions and pollution concentrations are measured, and polluters pay on the basis of how their emissions translate into increased pollution concentrations. As a result, the marginal costs of reducing pollution concentrations at the receptors are equalized among polluters, achieving a cost-effective reduction in pollution concentration.

This discussion relates directly to the fifth dimension of Professor Richards’ taxonomy, which concerns how polluters’ compliance with environmental regulation is measured.²⁰ The classic distinction, as Richards points out, is that between regulating output of pollution di-

18. For a discussion of ambient permit systems in particular, see generally Thomas H. TIETENBERG, *EMISSIONS TRADING: AN EXERCISE IN REFORMING POLLUTION POLICY* (1985).

19. See, e.g., Richard L. Revesz, *Federalism and Interstate Environmental Externalities*, 144 U. PENN. L. REV. 2341, 2410-2416 (1996).

20. See Richards, *supra* note 1, at 247.

rectly, and regulating the inputs to the production process that create the pollution. In both cases, the regulator is interested in the actual pollution output, rather than the inputs used; but the performance indicator differs. Professor Richards adopts a wider view, framing the issue not just in terms of inputs and outputs, but in terms of the “degree of correlation between the performance indicator and the desired outcome.”²¹ The choice of whether to regulate at the level of inputs or outputs, however, involves a tradeoff. As a general rule, we might suppose input controls, such as a tax on the sulfur content of coal, to be easier to administer than controls on emissions: the regulator needs only to track the coal burned by the power plant, rather than the more elusive flue gases escaping from the stacks. On the other hand, regulating inputs rather than pollution restricts the scope of the polluter’s discretion in limiting pollution. If “end-of-pipe” emissions reduction is cheaper than altering inputs, an input regulation will raise the costs of abatement.²²

We can extend this reasoning further by drawing on the concepts of ambient taxes or permits. Just as the appropriateness of an input-based regulation depends on the correlation between inputs and emissions, so the appropriateness of an emissions-based regulation depends on the correlation between emissions and environmental damages. In particular, the extent to which pollutant emissions from different sources are nonuniformly mixed—so that emissions from different sources have different impacts on the environment—determines how suitable an instrument focused on emissions is, compared to one based on ambient concentrations.²³

Ultimately, the “desired outcome” of pollution control policy is not a reduction in pollution emissions (or even aggregate emissions), any more than it is a reduction in inputs to polluting processes. Rather, the desired outcome is a reduction in *environmental damages*, whether that implies damages to ecosystems (e.g., acidification of Adirondack lakes) or increased health risks (e.g., from increased con-

21. *See id.*

22. The argument is essentially the same as in the case of other arbitrary restrictions on control techniques, such as technology standards. Prior to 1990, sulfur dioxide regulation was a classic example of a regulation resulting in higher costs than necessary. By effectively requiring power plants to install scrubbers (more precisely, mandating a certain percentage reduction in emissions, which could only be met with a scrubber), regardless of the sulfur dioxide emissions from the plant, the 1977 Clean Air Act Amendments limited the discretion of power plant operators and raised the costs of abatement. *See generally* BRUCE A. ACKERMAN & WILLIAM T. HASSLER, *CLEAN COAL/DIRTY AIR* (1981).

23. *See* TIETENBERG, *supra* note 18, at 22.

centrations of sulfur dioxide in urban areas). While emissions are closer to environmental damages than inputs, they are still somewhat removed. Thus, we can conceive of a series of possible levels of regulation, from inputs, to ultimate environmental damages. For a pollutant that causes human health damages, for example, this chain would run from inputs, to emissions, to ambient concentrations, to exposure, to human health risk. Indeed, we can also extend the argument in the other direction, viewing technology standards as one step prior to input-based regulations. Each step along the spectrum from technology to environmental damages represents a possible *locus of regulation*.²⁴

In some respects, the “ideal” instrument would be one that sets environmental damage as the locus of regulation. One can imagine a tax based on units of habitat loss, on fish killed from water pollution, or on the increased risk of morbidity from increased air pollution. Of course, administering such a tax would present insurmountable obstacles—notably, the problem of how to connect a given environmental effect with the polluter ultimately responsible. Put another way, the administrative costs of such a tax would outweigh the benefits from taxing the habitat loss or morbidity risk directly. Conceptually, however, such a tax merely occupies a different place on the same continuum along which we compared instruments based on pollution and those based on inputs. The fundamental tradeoff remains the same: the degree of correlation between the performance indicator and the desired outcome (lessened environmental damage), weighed against the administrative costs of the performance indicator chosen. The optimal location along the continuum—the locus of regulation—will vary with the context.

For example, it is feasible to monitor (and thus to regulate) sulfur dioxide emissions from the smokestacks of power plants. On the other hand, most proposals for global climate change policy involve taxes on carbon content in fuels or tradeable carbon rights, rather than taxes or permits denominated in carbon dioxide emissions. Such an approach is partly justified by the close correlation between carbon dioxide emissions and carbon content. At least as important, however, are the enormous administrative costs that would be involved in monitoring and taxing individual users of fossil fuels (e.g., automobile owners).²⁵

24. The author credits Robert Stavins with suggesting this framework.

25. To see the importance of the administrative costs involved in regulating emissions, suppose that a new “end-of-pipe” (i.e., post-combustion) technology to reduce carbon dioxide

The consideration of ambient permits and taxes, therefore, demonstrates that Professor Richards' discussion of the "degree of correlation between the performance indicator and the desired outcome" can be extended far beyond inputs and outputs. In fact, the same fundamental point applies to a whole spectrum of possible levels of regulatory intervention, from technology, to the underlying environmental damages. Indeed, a comprehensive taxonomy would include not only ambient permits and taxes, but also instruments tied directly to increases in human health risk or environmental deterioration.

V. HOW APPLICABLE IS TRANSACTION COST ECONOMICS TO ENVIRONMENTAL POLICY?

Professor Richards' most provocative argument is that transaction cost economics offers important lessons for instrument choice in environmental policy.²⁶ After reviewing the theory of contractual arrangements laid out largely by Oliver Williamson,²⁷ Professor Richards concludes that this theory is relevant to the choice of environmental policy instruments.²⁸ In particular, he argues that instruments (such as taxes or tradeable permits) that offer private parties more discretion over investments in pollution control also give rise to the possibility of regulatory opportunism. Fears that regulators might expropriate the gains from cost-reducing investments (e.g., by confiscating tradeable permits) may lead firms to under-invest in abatement activities.²⁹ Richards describes such distortions in firms' decision-making as "governance costs."³⁰ Since these costs are generated by allocating discretion to private parties, he argues, they provide potential grounds for preferring direct command-and-control regulation to market-based instruments.³¹

For all the originality of his argument, however, Professor Richards does not fully explain *why* transaction cost economics is rele-

emissions became widely available. Such a technology would diminish the correlation between carbon content and carbon dioxide emissions, since emissions could be reduced either by burning a fuel with lower carbon content or by installing the end-of-pipe abatement technology. However, the enormous administrative difficulties associated with regulating at the level of individual carbon dioxide emitters would remain, preserving a preference for regulating carbon content rather than carbon dioxide emissions.

26. See Richards, *supra* note 1, at 254-73.

27. See OLIVER E. WILLIAMSON, *THE ECONOMIC INSTITUTIONS OF CAPITALISM* (1985).

28. See Richards, *supra* note 1, at 261-65.

29. See *id.*

30. See *id.* at 258.

31. See *id.* at 265, 278-80.

vant to environmental policy. Rather, he simply asserts that the “factors that are key to the analysis of transaction costs—bounded rationality and opportunism on the part of the parties and asset specificity and uncertainty in the transaction—are equally relevant to pollution abatement programs.”³² Moreover, he implicitly embraces the premise that those characteristics are, in themselves, *sufficient* for the wholesale application of transaction cost economics to environmental policy.

In order to assess the applicability of transaction cost economics to environmental policy, I first review the basic structure of the “hold-up” problem at the heart of transaction cost theories, which motivates the design of hierarchical arrangements or the long-term contracts that mimic them. I then ask how well the model fits pollution control policy. In many cases, such as the examples provided by Professor Richards,³³ there may be some potential for expropriation by regulators, and resulting underinvestment by firms; but this problem is better conceived in terms of cost revelation by firms and time-inconsistency on the part of the regulator than as a contracting problem amenable to transaction cost analysis. In particular, the root of the commitment problem in environmental policy is the *regulatory relationship* itself, rather than the existence of an investment in a specific asset.

I next propose a particular version of a conventional “hold-up” problem that does apply to the context of environmental policy. Rather than being characteristic of market-based environmental regulation, however, this case is the result of a special assumption about regulatory behavior. In general, by instituting arms-length relationships between firms and the regulator, maintaining an essential anonymity among polluters, market-based instruments may tend to mitigate, rather than create, the hold-up problem. Further, even if we grant Professor Richards’ claim that a hold-up problem is likely to result from market-based instruments, his proposed remedy of a form of direct regulation such as a technology standard does not represent a solution.

Throughout my analysis, I will consider Professor Richards’ arguments within the context of a generic pollution abatement problem, in which emissions from a polluting industry are regulated by government. Because the instrument of “government production” is

32. *See id.* at 261-62.

33. *See id.* at 262-63.

somewhat elusive in this context (unless one takes it to mean state ownership of the industrial sector),³⁴ I will focus on what Professor Richards calls “direct regulation,” or technology standards—the form of “hierarchical control” suited to pollution abatement regulation.³⁵

A. *A Canonical Hold-Up Problem*

The question motivating transaction cost theories of the firm is, in some sense, why firms (hierarchical organizations characterized by vertical control among agents) exist at all. Given the well-known theoretical efficiency of perfectly functioning markets, why are spot market transactions insufficient? The answer, as advanced by Oliver Williamson in particular,³⁶ is that three characteristics of interactions (between firms or between agents) raise obstacles to the smooth functioning of contracts: bounded rationality, opportunism, and asset specificity. When agents are not fully rational, and lack perfect foresight, they cannot anticipate (and thus include in a contract) every possible eventuality that will affect their relationship. When agents are opportunistic, they seek to further their own self-interest, even at the expense of the other party. And when agents invest in assets that are specialized to the relationship, they risk being taken advantage of *ex post*—being “held up”—by the other party.

Consider the following stylized version of the hold-up problem.³⁷ Two firms, A and B, engage in a series of spot market transactions. Firm A supplies widgets to Firm B. Suppose that Firm A can install a specialized widget factory tailor-made to Firm B’s needs, at some extra cost. Suppose further that such an investment would be efficient: the gains from specialization would outweigh the costs of the investment. Under a system of spot-market transactions, Firm A has little

34. Indeed, one is never quite sure what Professor Richards has in mind by “government production” in most environmental policy contexts. His example—the National Forest system—suggests that what Professor Richards has in mind is “government production of a public good” (setting aside the issue of whether, in fact, the National Forests should be seen as a public environmental good, or as a storehouse of timber for private firms). But, the government cannot feasibly produce the good “clean air” (at least, after-the-fact removal of emissions from the air would, presumably, be vastly more expensive than end-of-pipe controls by emitters). The instrument of “government production,” therefore, appears to have little relevance for a large body of environmental policy.

35. See Richards, *supra* note 1, at 264.

36. See WILLIAMSON, *supra* note 27.

37. For an excellent and succinct review of transactions cost theories (and related theories of the firm), see OLIVER HART, *FIRMS, CONTRACTS, AND FINANCIAL STRUCTURE* 21-28 (1995). Hart also provides his own theories about vertical control and contractual relationships, based on a theory of complementary assets and residual rights in a firm’s production.

incentive to make the investment, because it will not be able to recoup the extra investment *ex post*. Once Firm A makes the investment, Firm B can exploit the bilateral relationship to extract the rents from the investment for itself. Having specialized to produce widgets for Firm B, Firm A's outside options are worse, because A's specific investment is of no use for the general widget market. In such cases, vertical relationships (or, at the least, contractual arrangements incorporating some features of hierarchical control) may solve the problem by aligning the incentives of the investing firm with its partner firm.

A prominent example of a potential hold-up problem and the vertical control used as a solution is the relationship between General Motors and Fisher Body, a long-time supplier of car bodies to the automobile manufacturing giant.³⁸ As Hart describes the relationship, it was subject to a hold-up problem: if either firm made an investment specialized to the relationship, it ran the risk that the other firm could expropriate some of the gains from that investment *ex post*, without having borne the costs. For example, if Fisher made an investment that allowed it to meet the specifications for GM auto bodies at a lower cost, it would not be able to recoup all of the cost savings of that investment. *Ex post* bargaining would allow GM to negotiate a lower price for auto bodies, taking advantage of Fisher's investment, which by definition did not yield cost savings in the production of other types of auto bodies. Thus, Fisher would have an inefficiently low incentive to invest.

Relationships between coal mines and coal-fired power plants provide further examples of specific investments: for example, a utility can locate a power plant adjacent to a coal mine, or install boilers dedicated to a certain type of coal. As transaction cost theories would predict, the duration of coal delivery contracts—a means of avoiding repeated bargaining and, thus, mitigating potential hold-up problems—is related to the degree of asset specificity.³⁹

It is important to stress that the particular *type* of hold-up problem at issue here grows out of the relationship created by the potential for asset specificity. The nature of the relationship between the two agents is altered critically by the specific investment, a process that Williamson labels the “fundamental transformation.”⁴⁰ A sup-

38. See *id.* at 6-8, 23-27.

39. See Paul L. Joskow, *Contract Duration and Relationship-Specific Investments: Empirical Evidence from Coal Markets*, 77 AM. ECON. REV. 168 (1987).

40. See WILLIAMSON, *supra* note 27, at 61.

plier, for example, may originally be determined by competitive bidding from among a number of candidates who are *ex ante* indistinguishable; but if the winning bidder then invests in specialized assets, it gains a distinct relationship with the contractor. In Williamson's terms, a "faceless" transaction in the marketplace (others use the phrase "arms-length relationship") has been replaced by a bilateral situation in which the identities of buyer and supplier matter.⁴¹ Rather than being ancillary to the relationship, the specialized investment is the linchpin; its existence creates the hold-up problem and the corresponding need for vertical integration or hierarchical control.⁴²

B. *Cost Revelation and the Regulatory Relationship*

To illustrate the hold-up problem in environmental regulation, Professor Richards cites two examples. On closer examination, however, neither quite represents the type of hold-up problem envisioned by Williamson and his colleagues. The first example is the EPA's emissions trading program in the 1980s, in which "fear of expropriation [of emissions permits by state agencies] likely impeded trading and, by implication, cost-effective investment."⁴³ In this instance, however, the fear of expropriation was based not on the regulator extracting the gains from a specific investment, but rather on the regulator learning new information from trading behavior and using it to ratchet up the level of regulation. That is, the key issue was not whether firms would make new investments to reduce their control costs, but whether they would reveal their existing control costs by trading emissions credits. While the reduction in trading may have reduced incentives for new investment, this was only an *indirect* effect of the fear of regulatory opportunism. In contrast, investment decisions are fundamental to the "hold-up problem" of transaction cost economics.

The second example cited by Professor Richards is that of the EPA's emissions standards for automobiles. Citing work by Yao, Professor Richards states that car manufacturers may underinvest in emissions control technologies in order to boost the cost of compliance and derail proposed emissions standards.⁴⁴ This case, however, again concerns revelation of cost information: car manufacturers

41. *See id.* at 61-62.

42. *See id.* at 61-63.

43. *See* Richards, *supra* note 1, at 263.

44. *See id.*

hope to inflate their costs so that the agency will choose a less stringent standard. Because the standards in question are announced in advance, any underinvestment by the automakers cannot be the product of fears that, *once the automakers invest in emissions control technology*, the gains they might otherwise get from such investments will be extracted by the regulator in the form of tighter standards.⁴⁵

Rather than representing contracting problems due to asset specificity and the resulting bilateral relationships, these two examples illustrate a different type of commitment problem, which I will label a *cost revelation problem*: the regulator is unable to commit in advance not to take advantage of cost information revealed by the firm's actions. If the firms reveal their cost of compliance to be low, for example, the regulator may have an incentive *ex post* to "ratchet up" the stringency of regulation. Anticipating this, the regulated firms have a strategic incentive to distort their own behavior in order to fool the regulator. This type of problem has been studied extensively in generalized regulatory settings, such as government procurement and cost regulation of natural monopolies.⁴⁶

The key distinction between this cost revelation problem and the contracting problem envisioned in the transaction cost literature is the nature of the relationship between the two parties. As described above, the core of the contracting problem in the transaction cost literature is the existence of a specialized investment, which, in effect, creates a new relationship out of a "faceless transaction." On the other hand, the cost revelation problem takes place within a relationship defined by the fact of regulation. The key element of the relationship is not the specific investment a firm makes, but the authority relationship defined by regulatory policy: the regulator, backed by the State, has the authority to require the firm to carry out actions (such as reducing its pollution) that the firm would not otherwise do.

45. There is necessarily a distinction between an investment that a firm would make with the expectation of a future payoff and one that would simply reveal cost information about the firm. In the case of EPA emissions standards, a firm that made an investment after the announcement of the stricter standards might do so to comply with the new rules, but it would not be hoping to capture savings in the costs of compliance with the old rules. The causal motivation behind underinvestment is avoidance of an indication to the EPA that the new standard was feasible; therefore, the firm might seek to fool the agency by not investing. The effect—underinvestment—is the same as in the hold-up problem, but this does not imply that the cause is also the same.

46. See generally JEAN-JACQUES LAFFONT & JEAN TIROLE, A THEORY OF INCENTIVES IN REGULATION AND PROCUREMENT (1993) [hereinafter LAFFONT & TIROLE, A THEORY]; Jean-Jacques Laffont & Jean Tirole, *Using Cost Observations to Regulate Firms*, 94 J. POL. ECON. 614 (1986).

Thus, the cost revelation problem applies equally well to the cost regulation of electric utilities,⁴⁷ where no specific asset is involved, as to the EPA's decisions on automobile standards. Firms in an environmental regulatory relationship may well make specialized investments, in the sense of pollution control investments that are not otherwise productive; but such investments are not the source of the special relationship between those firms and the regulator. Moreover, it is the revelation of information, rather than the decision to make a specific investment, that makes the firm vulnerable to regulatory opportunism. In the case of the emissions trading program, for example, I have argued that the question of investment is ancillary to the more basic issue of firm trading decisions. It is the trading decisions that will reveal information and are thus distorted by the regulatory relationship. The effect on investment is only secondary.

C. *Expropriation and Environmental Policy*

Professor Richards' examples may provide better illustrations of the cost-revelation problem than the hold-up problem; nonetheless, we can outline a hypothetical example that may capture what he has in mind. We can then use this example to assess Professor Richards' conclusions about governance costs and environmental policy instruments.

Suppose a regulator institutes a system of tradeable emissions permits for some specific air pollutant. To help distinguish the hold-up problem from the cost-revelation problem discussed above, suppose that the regulator has accurate information on polluters' control costs and can monitor any investments in new control technologies. Consider a polluting firm in the regulated industry contemplating an investment in a new abatement technology with a lower cost than its initial technology. Switching to the new technology will allow the firm to abate more cheaply and to sell the resulting excess emissions permits (or to reduce its purchases of such permits, if it is a net buyer).⁴⁸ If the benefit from doing so is greater than the cost of the investment, a cost-minimizing firm will adopt the technology.

Once firms have made their investments, however, the regulator can improve air quality by confiscating the excess permits of any firm who adopted the new technology. By doing so, the regulator would

47. See, e.g., Paul L. Joskow & Richard Schmalensee, *Incentive Regulation for Electric Utilities*, 4 YALE J. OF REG. 1 (1986).

48. This is where the theory of the dynamic incentive effects of regulation, discussed earlier, comes into play. See *supra* text accompanying notes 8 and 9.

extract the gains from the investment, in the form of improved environmental quality, at the expense of the firm. If polluters anticipate this possibility *ex ante*, they will discount the gains from their investments and thus have a lower incentive to adopt new technologies. Therefore, the regulator's ability to respond to new investment by tightening the stringency of regulation creates a hold-up problem.

An important aspect of this example, however, is that the basic result—that firms will underinvest in pollution control technology if they fear confiscation of the permits—is not robust with respect to changes in the scenario. The example above assumes that the regulator can, for each individual firm, determine which permits were created by the investment in technology and confiscate only those permits. Such specific action on the part of the regulator is necessary if the regulator is to target those cost savings that are due to the investment made by the firm. Suppose, on the contrary, that in response to investment by firms in the industry, the regulator lowers the *total* number of permits available—perhaps by lowering the total available at auction or by reducing permit allocations to all firms by some fixed proportion. Will the effect on firm investment be the same?

The answer, of course, is no. To take a simple case, assume that the firms are each “small enough” that an individual firm's decision to adopt a technology has only a negligible effect on the industry-wide marginal cost curve. In that case, the possibility that the regulator will reduce the number of emissions permits will actually *increase* the incentive to invest in new technology. This is because the cost savings from such investment are positively related to the price of permits: the more expensive it is to pollute, the greater the incentive to reduce abatement costs. Moreover, for any given distribution of marginal control costs among firms, the permit price will be higher when there are fewer permits available, simply as a result of supply and demand in the permit market.⁴⁹ Some firms that would not choose to adopt a technology under the original number of permits will therefore have an incentive to do so given the smaller number of permits, because the cost savings will be higher given the higher permit price. Thus, investment would *increase*, rather than decrease, as a result of the

49. If marginal abatement costs in the polluting industry fall due to technology adoption, then the permit price after adoption may be lower than the permit price before adoption, even if the regulator reduces the number of permits. In fact, we should expect such a result if the regulator is setting the number of permits optimally (taking into account the marginal damages function). However, *given* the adoption decisions by firms, the permit price will be higher if the regulator reduces the number of permits than it would be if the number of permits were fixed. Thus, reduction in the number of permits will increase the firm-level investment incentive.

regulator's inability to commit to a constant level of permits.⁵⁰ In more general settings, the regulator's inability to commit to a certain level of permits, and the *ex post* incentive to increase the stringency of regulation, may result either in over- or underinvestment, relative to the commitment case.⁵¹

Meanwhile, an emissions tax provides no obvious opening for such expropriation at all, at least not if we assume (with Professor Richards) that the regulator is motivated by economic efficiency. The optimal Pigouvian tax in any setting is equal to the marginal damage from pollution at the efficient level of pollution.⁵² If marginal damages are constant, the optimal tax is also constant, regardless of the polluting industry's marginal costs of abatement. If marginal damages are increasing per unit of pollution (i.e., marginal benefits of abatement are decreasing), a decrease in marginal abatement cost *decreases* the optimal tax. The scope for expropriation of the gains from investment is therefore limited, since the regulator's optimal response to investments in cost-reducing technology is either to keep the tax the same or to lower it. Moreover, the same result holds in the context of cost-effectiveness rather than efficiency—i.e., if the regulator seeks to achieve a given level of pollution control at least cost. The tax required to achieve a given target will fall with the marginal control costs.

Therefore, the hold-up problem described above depends crucially on both the instrument chosen (emissions permits rather than an emissions tax) and the manner in which the regulator employs it (confiscating the permits of specific firms, rather than reducing the total number of permits available). For the analogy with the transaction cost literature to hold, the regulator must confiscate the permits in a *targeted* fashion. The regulator must distinguish among polluters and treat them differently according to their investment decisions. In other words, the relationship between each polluting firm and the

50. In a sense, the firms face a collective action problem: as a whole, and the industry would be better off with more permits rather than fewer—but any individual firm, with a negligible effect on the regulator's decision, is better off adopting the technology than not. *See generally* Nathaniel O. Keohane, Policy Instruments and the Diffusion of Pollution Abatement Technology (1999) (unpublished manuscript, Harvard University) (on file with author).

51. *See generally* Gary Biglaiser, et. al., *Dynamic Pollution Regulation*, 8 J. REG. ECON. 33 (1995); Jean-Jacques Laffont and Jean Tirole, *Pollution Permits and Compliance Strategies*, 62 J. PUB. ECON. 85 (1996).

52. *See* BAUMOL & OATES, *supra* note 14, at 45. The notion of a corrective tax on pollution is generally credited to Pigou. *See generally* ARTHUR C. PIGOU, THE ECONOMICS OF WELFARE (1920).

regulator is determined by the investment decisions of that particular firm.

The guiding principle of market-based instruments such as emissions taxes or tradeable permits, however, is that they harness the power of the market to allocate pollution control responsibilities efficiently. Their crucial advantage derives from their institution of “arms-length” relationships between firms and the regulator. No firm has specific pollution abatement responsibilities assigned to it by the regulator. Instead, each firm makes its own choices based on a price signal (a tax or permit price) applicable to all polluters. The regulator makes no distinctions among firms, instead charging them all equally per unit of pollution or ensuring that their emissions do not exceed their permit holdings.⁵³

A hold-up problem, by contrast, requires that the regulator’s treatment of polluters depend not just on their emissions, but also on their investment decisions. Thus, the arms-length relationships characteristic of market-based instruments are at odds with the necessary conditions for the hold-up problem. Of course, as our hypothetical example suggests, one can imagine regulatory opportunism in combination with market-based instruments. However, such opportunism is hardly integral to such instruments, as Professor Richards seems to suggest.⁵⁴ If anything, the regulator–firm relationship required for the hold-up problem is fundamentally *antithetical* to the decentralized structure of market-based instruments.

D. *Would Direct Regulation Solve the Problem?*

Professor Richards, having argued that a hold-up problem exists with market-based instruments, proposes that direct regulation—in particular, technology standards—could solve the problem. In the preceding section, we showed that only under special assumptions will a market-based instrument create the hold-up problem he envisions. In this section, we ask whether technology standards would fare any better, *even granting the existence of a hold-up problem for market-based instruments*.

Under a pure technology standard, the regulator requires each polluting firm to install a particular technology. As Professor Ri-

53. In the case of ambient taxes or permits, discussed *supra* Part IV, the regulator does distinguish among firms based on their locations. But the principle of “arms-length” relationships remains, since the rule used to distinguish among firms is a general one, and (once a firm has chosen a location) does not depend on a firm’s actions.

54. See Richards, *supra* note 1, at 265, 278-80.

chards argues, this is tantamount to hierarchical government control over pollution abatement decisions. Because hierarchical arrangements have been suggested as solutions to the contracting problems studied in the transaction cost literature, Professor Richards concludes that they are also the solution in the environmental policy arena.⁵⁵ However, he never spells out *how* direct control of abatement technology would alleviate the hold-up problem he perceives.⁵⁶ We consider three possible arguments in turn.

One argument might be that under a technology standard, a firm does not earn cost savings from adopting a new technology, since there is a zero price on emissions. Moreover, the regulator obviously cannot expropriate cost savings that do not exist. Nonetheless, the proper solution by the regulator to rent extraction can hardly be to eliminate those rents altogether. To remove the distortionary effects on firms' decisions by preventing them from making those decisions may dispose of the hold-up problem, narrowly defined, but it does not advance the underlying goal of encouraging efficient investment decisions. What would be the reason, in this case, to suppose that the regulator's choice of technology would be any better than the firm's choice, even given the distortionary effects of the hold-up problem? The usual tradeoff between "high-powered incentives" and rent extraction in regulation, which Professor Richards alludes to, does not bear on the choice between market-based instruments and direct regulation.⁵⁷ The former tradeoff involves the choice between the regulator's desire to extract rents (e.g., to prevent "windfall profits") and the regulator's desire to create incentives for cost-reducing innovation by the firm.⁵⁸ These objectives conflict with one another, because the prospect of rents provides the firm with the incentive to reduce costs. In the case of pollution control, however, high-powered incentives and the potential for rent extraction are both products of market-based instruments. Direct regulation would eliminate the opportunity for rent extraction, but it would also eliminate the incentives for cost reduction.

55. See Richards, *supra* note 1, at 265.

56. Similarly, the transaction cost literature has itself been criticized for being vague on precisely *why* vertical integration should solve the hold-up problem. See HART, *supra* note 37, at 27.

57. See Richards, *supra* note 1, at 267.

58. See generally LAFFONT & TIROLE, A THEORY, *supra* note 47; Joskow & Schmalensee, *supra* note 46.

A second justification for technology standards relates to “adaptability.” Professor Richards argues that “the government may want to employ government production or direct regulation if it anticipates the need to adapt its environmental programs to changes in information or changes in goals.”⁵⁹ But Professor Richards fails to explain why a technology standard (a form of “direct regulation”) would be more flexible than a tax or a tradeable permit system. A technology standard for new sources might well be flexible, in the sense that a change from requiring technology A to be installed in new sources to mandating technology B can be made with negligible impact on existing polluters. A change in a tax or tradeable permit system, on the other hand, will presumably affect all polluters, not just new sources, and therefore has a greater effect on the technology decisions of existing firms.

However, this type of flexibility does not imply a corresponding flexibility in response to changes in the regulator’s *goals*. Those goals, after all, are defined not by the type of technology being installed in new sources, but by the overall level of pollution. Professor Richards’ original cost-minimization problem, after all, includes the “pollution abatement requirement” as a key constraint. The natural extension of a pollution abatement requirement to a setting in which new sources appear is a target level of overall pollution emissions.⁶⁰ The policy goal is thus reasonably assumed to be a level of pollution emissions, rather than a distribution of technologies. Indeed, the latter goal has no direct relevance for net social benefits, except to the extent that it influences the ultimate level of pollution.⁶¹

When we recall that the ultimate goal of pollution control policy is some target level of pollution, the apparent flexibility of technology standards vanishes. If the regulator wishes to reduce pollution, it will not be sufficient, in general, to change technology requirements for new sources: the technologies used by existing sources will have to change as well. Given that fact, then, technology standards offer no greater “adaptability” than other policy instruments. In fact, technology standards are *less* flexible than other instruments, because the

59. Richards, *supra* note 1, at 264.

60. That is, a fixed level of pollution abatement will result in greater overall pollution emissions if the number of sources of pollution grows. Of course, a regulator might well take such economic growth into account in determining the desired level of pollution at different points in time.

61. As discussed *supra* Part IV, the ultimate goal is presumably in terms of environmental quality. The distinction may be important when pollution from different sources is nonuniformly mixed.

choice of abatement technology is dictated by the regulator, rather than left to the discretion of individual firms.

A third possible argument for why technology standards might alleviate the hold-up problem is that hierarchical control arrangements tend to align the interests of the two parties in the hold-up relationship. This type of argument is a plausible justification for vertical control and hierarchical arrangements in the traditional setting of transaction cost economics. For example, in the case of GM and Fisher Body described above, the merger of the two companies brings the interests of the management of Fisher Body in line with the interests of GM, since the former is subordinated to the latter. As a result, the hold-up problem may be alleviated, since GM no longer needs to worry that Fisher Body may try to take advantage of GM's investments in using Fisher auto bodies; similarly, the gains to specialization by Fisher Body will now be captured by the entity which bears the costs of investment (that is, the merged firm).

However, it is difficult to see how such a benefit might emerge from technology standards. Even if the regulator dictates the pollution control technology to be used by a firm, the firm still bears the cost of buying that technology. The regulator thus lacks any incentive to ensure that the technology is cost-effective. Establishing this form of "hierarchical control" does not align the interests of the regulator with those of the polluting firm. Those interests would seem to be fully aligned only under state ownership of the polluting industry, which seems a rather extravagant solution to the hold-up problem—especially given the doubts raised earlier about whether the hold-up problem even exists in the environmental arena.

VI. CONCLUSION

Professor Richards has advanced a taxonomy of environmental policy instruments that ably demonstrates a range of possible connections among the various forms of regulatory policy—from direct regulation such as technology standards to market-based instruments such as emissions taxes. However, his effort to use the taxonomy as a basis for comparing and evaluating the instruments is subject to several criticisms.

First, he employs a static cost-minimization problem, rather than a dynamic one. Of course, Professor Richards touches on the dynamic effects of policy instruments throughout his article—in particular, dynamic effects are implicit in the discussion of investment decisions and regulatory opportunism. Explicitly incorporating the

time dimension into the framing of the regulator's objective problem would give his taxonomy a more coherent logical structure. Moreover, a focus on the time dimension suggests a few new points, such as the possibility that political and legal constraints evolve over time, possibly as a function of the "stock" of accumulated knowledge and experience.

Second, Professor Richards blurs the distinction between emissions payments—transfers from polluters to the regulator under instruments such as auctioned tradeable permits and emissions taxes—and the environmental damages from pollution. Emissions payments need not be equal to environmental damages. More fundamentally, emissions payments are best viewed as a means of achieving efficiency, rather than distributional equity. Indeed, compensation to victims of pollution is outside the scope of the efficient solution to the pollution externality, and—like any other form of government redistribution—comes at some cost in terms of deadweight loss. Moreover, victim compensation is incompatible with "revenue recycling," which is necessary to minimize the distortionary effects of environmental regulation in the presence of pre-existing tax distortions.

Third, Professor Richards overlooks an important class of policy instruments: ambient taxes and permits, which attach a price directly to increased pollution concentrations, rather than to the precursor of those concentrations—pollution emissions. The relationship between ambient permits or taxes, and emissions permits or taxes, is exactly parallel to that between an emissions standard and an input standard. In both cases, the instruments being compared occupy different points along the spectrum of the "locus of regulation." The ultimate goal of environmental policy is to reduce environmental damages, whether in the form of deterioration of environmental quality or in the form of increased human health risk. In theory, policy instruments could be designed to be contingent upon increased cancer risk in a population, resulting, say, from a given firm's pollution emissions. Such an instrument is not used in practice not because it is inconceivable, but because there are enormous practical obstacles to implementing it—foremost among them the difficulty of connecting the emissions of a given source to the actual increased health risk in the population.

The choice between regulating emissions and regulating contributions to ambient concentrations involves weighing the benefits of the latter approach (a closer correlation between the object of regulation and the environmental outcome) against its costs (the adminis-

trative difficulties inherent in attributing ambient concentrations to specific sources). The same tradeoff applies to the choice between regulating emissions and regulating inputs, as in the case of carbon dioxide. Thus, the choice between using inputs or outputs as the “performance indicator” for assessing compliance with environmental regulation can be viewed within the larger context of choosing the locus of regulation: that decision of where regulatory intervention should be established along the chain from the control technology and the inputs used in production all the way to environmental impacts.

Fourth, Professor Richards attempts to apply the lessons of transaction cost economics to the realm of environmental policy. The examples of hold-up problems that he offers, however, are better seen as examples of cost-revelation problems in which the regulated firms distort their behavior in order to hide information about their true costs from the regulator. The relationship between the government and the firm in environmental regulation is defined by the fact of regulation, rather than by the existence of a specialized asset (as in the transaction cost literature). While we can conceive of a hold-up problem in environmental regulation, it requires the special assumption that the regulator distinguishes among firms on the basis of their investment. Such an arrangement runs counter to the spirit of market-based instruments, which preserve a kind of anonymity among regulated firms. Thus, the very nature of market-based instruments will tend to mitigate the hold-up problem. Finally, even if we grant the existence of a hold-up problem under some forms of incentive-based regulation, there is no reason to suppose that a technology standard will fare any better. On balance, therefore, transaction cost economics is only of limited applicability to environmental policy. More importantly, it does not provide a sound basis for choosing among policy instruments. “Governance costs” do not offer a reason to prefer command-and-control regulation to market-based instruments.

Ultimately, the primary contribution of Professor Richards’ article is to elucidate the relationships among various instruments and to point out some connections that have been overlooked, such as the dual relationship between subsidies and contracts mirroring that between emissions taxes and permits. Richards’ discussion, however, adds little to what we already know about choosing among policy instruments.