

NBER WORKING PAPER SERIES

SETTING THE INITIAL TIME-PROFILE OF CLIMATE POLICY:
THE ECONOMICS OF ENVIRONMENTAL POLICY PHASE-INS

Roberton C. Williams III

Working Paper 16120
<http://www.nber.org/papers/w16120>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2010

For their helpful comments and suggestions, I thank Don Fullerton, Stephen Holland, and participants in the NBER Conference on The Design and Implementation of U.S. Climate Policy. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2010 by Roberton C. Williams III. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Setting the Initial Time-Profile of Climate Policy: The Economics of Environmental Policy
Phase-Ins

Roberton C. Williams III
NBER Working Paper No. 16120
June 2010
JEL No. D62,H23,Q52,Q54,Q58

ABSTRACT

This paper considers the question of under what circumstances a new environmental regulation should “phase in” gradually over time, rather than being immediately implemented at full force. The paper focuses particularly on climate policy, though its insights are more general. It shows that while adjustment costs provide a strong efficiency argument for phasing in a quantity-based regulation (or allowing intertemporal flexibility that creates the equivalent of a phase-in), this argument does not apply for price-based regulation. Indeed, in many cases, it will be more efficient to do just the opposite, setting an initially very high emissions price that then falls as the policy phases in. This difference in results comes not from any fundamental difference between price and quantity policies: under either policy, the efficient quantity of abatement rises over time, while the efficient price stays constant or even falls. But other considerations, such as distributional concerns or monitoring and enforcement issues, may still argue for a gradual phase-in even for a price-based policy.

Roberton C. Williams III
Department of Agricultural and Resource Economics
University of Maryland,
Symons Hall
College Park, MD 20742
and NBER
rwilliams@arec.umd.edu

This paper considers the question of under what circumstances a new environmental regulation should “phase in” gradually over time, starting with an initially lax regulation and then gradually tightening, rather than being immediately implemented at full force. Phase-ins are a very common – perhaps ubiquitous – feature of new environmental regulations, and can greatly influence the near-term costs and benefits of policy.

This differs from the broader, longer-term question of whether regulation should tighten over time. There are other reasons why it might be efficient for regulation to gradually tighten (e.g., if incomes are rising and thus the willingness to pay for a cleaner environment is also rising). The key distinction is that with a phase-in, the reason for gradually tightening over time is because the policy is new. A natural argument for a phase-in is that it provides time for individuals and firms to adjust to the new policy. Therefore, much of this paper focuses on the role of adjustment costs.

Prior work on the broader issue of the optimal time-profile of climate policy has indirectly addressed the issue of phase-ins. For example, Wigley *et al.* (1996) shows that because of capital adjustment costs, the least-cost path to achieve a given atmospheric concentration of CO₂ departs only gradually from the business-as-usual path, thus implicitly suggesting some sort of phased-in policy. And a substantial literature focuses on the question of whether learning-by-doing accelerates or slows the optimal pace of carbon abatement (e.g., Goulder and Mathai, 2000, or Manne and Richels, 2004), a question that implicitly relates to phase-ins. However, none of these papers specifically considers the phase-in question, or separates out this question from other influences on the optimal abatement path. And to my knowledge, no prior work even implicitly addresses the phase-in question in a general context, or in any specific environmental context other than carbon abatement, even though phase-ins have

been included in many other environmental regulations.¹

This paper uses an analytical dynamic model to consider the phase-in question in a general environmental regulation context, and then discusses implications of that model in the specific context of climate policy. The paper shows that while adjustment costs provide a strong efficiency argument for phasing in a quantity-based regulation (or allowing intertemporal flexibility that creates the equivalent of a phase-in), this argument does not apply for price-based regulation. Indeed, in many cases, it will be more efficient to do just the opposite, setting an initially very high emissions price that then gradually falls over time. This difference in results comes not from any fundamental difference between price and quantity policies, but simply from a difference in how one defines whether the policy is phased-in or not: under either policy, the efficient quantity of abatement rises over time, while the efficient price stays constant or even falls. However, other considerations, such as distributional concerns or monitoring and enforcement issues, may still argue for a gradual phase-in even for a price-based policy.

The next section of this paper presents a simple analytical dynamic model of environmental regulation and uses that model to address the phase-in question. The following section considers possible extensions to that model that might provide a further rationale for a gradual phase-in of a new regulation. A final section concludes and discusses implications for policy.

I. A Simple Model

This section introduces a simple analytical dynamic model of environmental regulation, and uses that model to address the question of under what circumstances an environmental policy

¹ Montero (2000) addresses a different issue: the optimal design of a trading program that allows otherwise unregulated sources to opt-in to the program. Such opt-in provisions have been included in early phases of a number of emissions trading programs (most notably the U.S. SO₂ trading program).

should be phased in gradually rather than immediately implemented at full force. A key element of this problem is that capital cannot instantly adjust in response to policy. This provides the main argument for phasing in policy: a gradual phase-in avoids making existing capital prematurely obsolete, and allows time to build up a stock of less-polluting capital.

To incorporate this issue, production follows

$$(1) \quad Y_t = F(H_t, E_t),$$

where Y is output of a pollution-intensive good, H is the stock of pollution-intensive capital, and E is the pollution emissions rate. The production function is concave and twice-differentiable.

In addition, pollution and capital are complements, so $\frac{\partial^2 F}{\partial H \partial E} > 0$.

For simplicity, this model explicitly considers only a single type of capital, and production of only one good. This is probably best understood as a partial-equilibrium model, with the single good representing an aggregate of output from all pollution-intensive industries. A model with two distinct types of capital, one polluting and one non-polluting (or less-polluting), or with production of both pollution-intensive and non-pollution intensive goods, would be more complex but would yield fundamentally the same results.

Capital depreciates at the rate δ , and thus the rate of change of the capital stock is given by

$$(2) \quad \dot{H}_t = I_t - \delta H_t,$$

where I is the rate of investment (or disinvestment, if negative). The cost of investment is given by

$$(3) \quad C(I_t),$$

which is strictly convex and twice-differentiable. This function includes the cost of the capital

itself (which will be negative if I is negative), plus any adjustment cost. The profit-maximization problem of a representative firm is then given by

$$(4) \quad \max_{E, I} \int_0^{\infty} [p_t F(H_t, E_t) - C(I_t) - \tau_t E_t] e^{-rt} dt ,$$

subject to the capital transition equation (2), where p is the price of output, τ is the emissions tax rate or emissions permit price, and r is the discount rate. The first-order condition for the emissions rate is then

$$(5) \quad p_t \frac{\partial F}{\partial E_t} = \tau_t ,$$

which equates the marginal value product of emissions with the emissions tax rate. The first-order condition for investment is

$$(6) \quad \frac{\partial C}{\partial I_t} = \lambda_t ,$$

which sets the marginal cost of capital equal to its current-value shadow price, λ . The costate equation gives the rate of change of λ as

$$(7) \quad \dot{\lambda}_t = (r + \delta) \lambda_t - p \partial F / \partial H_t .$$

The intuition for this equation is that the return on capital (its marginal value product plus the change in its shadow price) must equal the cost of holding capital (the discount rate plus the depreciation rate, times the shadow price).

The next sub-section considers regulation of a flow pollutant (one for which pollution damage is caused entirely by the current flow of emissions). The following sub-section then considers regulation of a stock pollutant (one for which damage is caused by the accumulated stock of emissions, as is the case for greenhouse gases).

A. Regulation of a Flow Pollutant

In the flow pollutant case, pollution damage will be given by the function $D(E_t)$, which is increasing, convex, and twice differentiable. The regulator's problem is given by

$$(8) \quad \max_{\tau} \int_0^{\infty} [p_t F(H_t, E_t) - C(I_t) - D(E_t)] e^{-rt} dt ,$$

which is very similar to the firm's problem (4), except that in the regulator's objective the cost of pollution is the pollution damage done, whereas the analogous term in the firm's objective is the emissions tax paid. The regulator's first-order condition for the emissions tax rate is

$$(9) \quad p_t \frac{\partial F}{\partial E_t} = \partial D / \partial E_t ,$$

which sets the marginal value product of emissions equal to the marginal damage. The regulator can achieve this by setting the emissions tax rate equal to marginal damage,

$$(10) \quad \tau_t = \partial D / \partial E_t ,$$

which causes the firm's first-order condition (5) to be equivalent to the regulator's first-order condition (9). If the pollution tax is set equal to marginal damage at all points in time, then the firm's first-order condition for investment (6) and costate equation (7) will also be equivalent to the analogous equations for the regulator. Just as in a simple static model, the optimal emissions tax simply equals the marginal damage from emissions.

What does this imply for phase-ins? First, consider the case in which the marginal damage from pollution is constant (i.e., damage is linear in emissions). The optimal emissions tax, equal to marginal damage, will then also be constant. Thus, in this case, it is not optimal to phase in the emissions price: the optimal path has the emissions price go immediately to its fully-phased-in level and stay constant at that level.

However, the optimal time path for emissions in this case does involve a phase-in.

Imposing a constant emissions price causes an immediate drop in emissions. Because capital and emissions are complements, that drop in emissions causes a corresponding drop in the marginal product of capital, which in turn causes the shadow price of capital to fall, leading to a reduction in investment. That drop in investment means that the quantity of capital will gradually fall, with a corresponding gradual fall in the emissions rate (again, because capital and emissions are complements), eventually converging to a new steady state with lower levels of capital and emissions.²

Thus, the optimal policy doesn't phase in the emissions price, but does phase in the emissions quantity. Regardless of how quickly or slowly capital can adjust, setting the emissions price equal to marginal damage internalizes the externality and thus leads to the efficient level of emissions. This might lead to earlier retirement of polluting capital and to higher costs than would a gradual increase in the emissions price, but if so, then retiring that polluting capital earlier and incurring higher costs is efficient. But because the capital stock takes time to adjust, the level of emissions reductions implied by any given emissions price will rise over time, thus gradually phasing in the emissions reductions. Another way of thinking about this is to view it as a higher price elasticity of emissions in the long run than in the short run, so the same emissions price will lead to a greater emissions reduction in the long run than in the short run.

Now consider the case in which marginal damage is increasing in the level of emissions. As just shown, a constant emissions price implies a gradually falling level of emissions over time. But if marginal damage is increasing in emissions, a gradually falling level of emissions implies that marginal damage will also be gradually falling. Therefore the optimal path must

² Note that if emissions and capital were substitutes, as would be the case for abatement capital, the optimal path would still entail a gradual drop in emissions. The chain of reasoning is the same as for the complements case, except that the signs of the changes in the shadow price of capital, investment rate, and quantity of capital are all reversed. A similar logic would apply in a model with both polluting and non-polluting capital: along the optimal path, the level of polluting capital would fall and the level of non-polluting capital would rise.

entail an emissions price that initially jumps to a level above its long-run level and then gradually falls over time as the capital stock adjusts. In this case, not only does the optimal path not entail a phase-in of the emissions price, but it actually implies the opposite.

Going one step further, consider a case with threshold damages: marginal damage is very low up to some threshold level of emissions and very high beyond that threshold. In this case, the optimal policy will hold the level of emissions right at that threshold. This will require a very high initial emissions price that then gradually falls over time. Thus, in this case, the quantity of emissions jumps immediately to its long-run level, without any phase-in, while the emissions price phases-in in reverse, starting high and then falling over time.

B. Regulation of a Stock Pollutant

Now consider the case of a stock pollutant such as greenhouse gases. Let P_t represent the stock of pollution, and $D(P_t)$ the damage caused. In this case, the regulator's problem is

$$(11) \quad \max_{\tau} \int_0^{\infty} [p_t F(H_t, E_t) - C(I_t) - D(P_t)] e^{-rt} dt,$$

subject to the same capital transition equation (2) and to a transition equation for the pollution stock, given by

$$(12) \quad \dot{P}_t = E_t - \eta P_t,$$

where η is the natural rate of decay of the pollution stock. The regulator's first-order condition for the emissions price is now

$$(13) \quad p_t \frac{\partial F}{\partial E_t} = \mu_t,$$

where μ_t is the shadow price of emissions at time t . The costate equation for μ_t is

$$(14) \quad \dot{\mu}_t = (r + \eta) \mu_t - \partial D / \partial P_t,$$

which can be solved to give

$$(15) \quad \mu_t = \int_0^\infty \frac{\partial D}{\partial P_{t+i}} e^{-(r+\eta)i} di .$$

Just as in the flow pollutant case, the first-order condition equates the marginal benefit from emissions with the marginal damage from emissions, which in this case is the discounted value of future pollution damage caused by a marginal unit of emissions at time t .

As in the flow pollutant case, it is helpful first to consider the case in which marginal pollution damage is constant. In this case, μ_t will also be constant (as can be seen by examining (14) or (15)), which in turn implies a constant optimal emissions price. Just as in the flow pollutant case, the optimal path has the emissions price jump immediately to its long-run level and then stay constant, while emissions gradually fall over time – in other words, the emissions price is not phased-in, but the emissions quantity is.

For the case in which marginal pollution damage is increasing in the stock of pollution, the results are again similar to the analogous results for a flow pollutant. A constant emissions price would imply a gradual fall in emissions, which in this case implies a gradual fall in μ_t (again, this can be seen by examining (14) or (15)). Thus, the optimal path must entail an emissions price that initially jumps to a level above its long-run level and then gradually falls over time.³ This effect will be much less pronounced than it would be for a flow pollutant (because a gradual fall in $\partial D / \partial P_t$ implies a much slower fall in μ_t), but it nonetheless demonstrates the same pattern, which is the opposite of the usual phase-in.

³ This argument assumes that the stock of pollution is at a steady state prior to the introduction of any regulation. This is not the case for carbon dioxide, for which the stock of pollution is currently rising rapidly. In such a case, the shadow price of emissions will follow a path similar to that of the pollution stock, initially rising, possibly overshooting its long-run level, and then eventually converging to a steady state. This may resemble a phase-in, because of the initially rising optimal emissions price, but arises for different reasons.

II. Possible Alternative Justifications for Phase-Ins

The previous section's results show that capital adjustment costs imply that phasing in the quantity of emissions reductions is optimal, but that phasing in the emissions price is not: the optimal emissions price jumps immediately to a level at or above its long-term level, without any phase-in period. Nonetheless, many environmental regulations have gradually phased in both the quantity and price of emissions. Phase 1 of the U.S. SO₂ trading program covered only a small fraction of the pollution sources that were covered in Phase 2, so for those not covered by Phase 1, the emissions price they faced was clearly higher in Phase 2. And even for those sources covered during Phase 1, the emissions caps in Phase 2 were enough tighter to imply a higher emissions price. Similarly, under the EU Emissions Trading System for carbon, the second-phase caps were enough tighter than the caps during the first phase to imply a substantially higher permit price. Moreover, environmental regulations are almost always announced well before they are to take effect, which also represents a phase-in.

Were the initial phases of those programs inefficiently designed, or do other factors provide some justification for phasing in both the quantity and price of emissions? This section discusses two such extensions to the model: distributional concerns and monitoring and enforcement issues.

A. Distributional Considerations

The model in Section I assumed that the regulator is setting policy to maximize efficiency. In practice, however, distributional considerations are often at least as important as efficiency, and policy decisions frequently represent a compromise between these two factors.

Suppose that in addition to maximizing efficiency, the regulator would also like to limit the cost imposed on firms (or more generally, on the owners of pollution-intensive capital).

Under some circumstances, this additional goal could imply a gradual phase-in of the emissions price as well as the emissions quantity.⁴

Consider an extreme case as an illustrative example: suppose that emissions and polluting capital are perfect complements in production (i.e., the production function (1) is Leontief), polluting capital cannot be liquidated once it is installed (i.e., $C(I_t) \geq 0$ for $I_t < 0$), and the efficiency-maximizing emissions price is not high enough to cause capital to be idled, but is more than enough to stop new investment. Thus, on the optimal path, pollution-intensive production continues, but the stock of polluting capital is allowed to depreciate over time, eventually converging to zero.

In this case, announcing that an emissions tax will be imposed at some future date but not imposing any tax before that date can have the same effect on investment and emissions as immediately imposing a tax. If the future date is not too distant (sufficiently near to cause the shadow value of capital to drop below the marginal cost of investment), then in either case, investment will stop immediately, and emissions will fall gradually as the capital stock depreciates. Thus, the efficiency consequences of these two policies are identical, but waiting to impose the tax reduces the cost to capital owners. If the regulator puts more weight on the cost to capital owners than on government revenue, then waiting to impose the tax would be optimal.

In less extreme cases, phasing in the emissions price will have some efficiency cost, but that cost could still be outweighed by distributional considerations. However, in these cases, an emissions price phase-in would still be a second-best policy: the regulator could achieve the

⁴ This is similar to Feldstein's (1976) argument for announcing a tax reform well in advance of the date it will take effect, which is less efficient than having it take effect immediately but may still be worthwhile for distributional reasons. Zodrow (1985) considered this argument in a dynamic model with capital adjustment costs, and found that whether it justifies some form of phase-in depends on the magnitude of capital adjustment costs. The optimal capital tax problem that Feldstein and Zodrow consider differs substantially from the optimal environmental tax problem considered here, but nonetheless it seems likely that a similar result would hold here.

same distributional outcome at lower efficiency cost by immediately imposing an emissions price equal to marginal damage and providing a compensating transfer (which could take the form of emissions permit allocations or inframarginal exemptions from an emissions tax) to owners of polluting capital. Only if such transfers aren't possible would an emissions price phase-in be optimal.

B. Monitoring and Enforcement

The model in Section I also ignores issues of emissions monitoring and enforcement of regulations. Incorporating such issues might provide another argument for phase-ins. Suppose that the regulatory agency has a limited capacity for monitoring and enforcement, and that increasing that capacity will take time (one could view this as accumulating “enforcement capital”). In such a case, it could be optimal initially to regulate only a relatively small set of polluters, those who would be expected to achieve relatively large reductions in emissions at relatively low cost. Then the set of regulated firms could be expanded over time as the regulatory agency’s enforcement capacity grows. The resulting phase-in policy would look much like the phase-in of the U.S. SO₂ trading program, which started by regulating a relatively small number of large and particularly pollution-intensive plants in Phase 1, and then expanded to include smaller and less-polluting plants in Phase 2.

Again, though, it is not at all clear that such a policy is genuinely optimal. It might well be more efficient for the regulation to cover all polluters immediately, but for limited enforcement resources to be directed primarily (though not exclusively) toward the largest polluters.

III. Conclusions

This paper has shown that capital adjustment costs provide an efficiency justification for a gradual phase-in of the quantity of emissions reductions under a new environmental regulation. But this argument does not hold for phase-ins of the emissions price. Indeed, the optimal policy is just the opposite – the emissions price immediately jumps to a point above its long-run level, and then gradually declines to that long-run level over time – for any case in which marginal pollution damage is increasing in the quantity of emissions.

This result calls into question the approach taken with many environmental regulations, which have gradual phase-ins of both the quantity of emissions reductions and the emissions price. Given this paper's simple and highly stylized model, it certainly cannot rule out the possibility that there are other considerations that would justify such phase-ins, and further work to explore such possible justifications would be valuable. But these results do suggest that a more aggressive emissions price path in the initial implementation of a new regulation could be more efficient.

References

Feldstein, Martin, 1976. "On the Theory of Tax Reform." *Journal of Public Economics* 6:77-104.

Goulder, Lawrence, and Koshy Mathai, 2000, "Optimal CO₂ Abatement in the Presence of Induced Technological Change." *Journal of Environmental Economics and Management* 39:1-38.

Manne, Alan, and Richard Richels, 2004, "The Impact of Learning-by-Doing on the Timing and Costs of CO₂ Abatement." *Energy Economics* 26:603-619.

Montero, Juan-Pablo, 2000, "Optimal Design of a Phase-In Emissions Trading Program." *Journal of Public Economics* 75:273-291.

Wigley, T., R. Richels, and J. Edmonds, 1996, "Economic and Environmental Choices in the Stabilization of Atmospheric CO₂ Concentrations." *Nature* 379:240-243.

Zodrow, George, 1985. "Optimal Tax Reform in the Presence of Adjustment Costs." *Journal of Public Economics* 27:211-230.