The Strategic Benefits of Uniform Environmental Standards

Anthony Heyes (Royal Holloway, University of London)

and

Kenneth L. Simons (Rensselaer Polytechnic Institute)

Short title:

[1]Heyes (corresponding author) is at Department of Economics, Royal Holloway College, University of London, Egham, Surrey TW20 0EX, England (a.heyes@rhul.ac.uk). We are grateful to Peth Tuppe, Sandeep Kapur, Michael Crew, seminar participants at UCL and the World Bank, and two referees of this journal for very constructive comments.

Abstract. Regulators are sometimes prevented from setting standards on a firm-by-firm basis. Such restrictions seem inefficient, and rationales for their prevalance have been in terms of politics or fairness. While the requirement that regulation be uniform does not allow the agency to commit to future standards, it does mean that what is required from one firm must be required from them all. Regulated firms are strategic about what information they reveal to regulators - either directly or by their actions. Because the stringency of regulation faced by any particular firm is tied to industry average characteristics, that firm is less prone to hide its access to a low-cost compliance technology by inefficient choice of technique. As such the requirement of uniformity offers partial mitigation of an informational ratchet problem which makes regulation based on 'averages' attractive. We characterise alternative settings in which it is desirable or undesirable to require that standards be uniform. **Keywords: Environmental regulation - strategic interaction - ratchet principle**

1. Introduction

In many settings regulators are obliged to treat firms uniformly. In environmental regulation, this most commonly means an obligation to set uniform rather than firm-specific standards.

Given that firms are likely to vary in their compliance costs this seems inefficient - you would want firms with higher marginal compliance costs to be faced with less demanding requirements - and the rationales given for why regulatory agencies may be placed under such constraints have typically been framed in terms of politics or notions of 'fairness'. Kolstad (1987) provides a number of (non-economic) arguments for why uniform standards may be popular, noting that legislators often consider it "... inequitable and inappropriate to let regulators treat polluters differentially, or on a firm-by-firm basis" (Kolstad (1987: 591)).[1] Boyer and Laffont (1999) is an example of a paper that provides a political-economic rationale for a constitution requiring the use of (otherwise) inefficient policy instruments in terms of limitations that such instruments place on the ability of politicians to distribute rents.

But if we ignore these sorts of concerns can it ever be efficient to require uniformity? This is the question that we explore here.

---

[1] The meaning of equitable here is, of course, far from unambiguous. For miscellaneous environmental examples, and discussion of what alternative interpretations could be attached to 'equal treatment' see Lichtenberg, Spear and Zilberman (1993), Ashworth and Papps (1991).

The benefits of letting regulators 'tailor' requirements to fit the circumstances of individual firms are well-known - first-best standards can be set to equate the marginal costs and benefits of abatement at the firm-level - and conventional wisdom asserts that preventing tailoring when firms are heterogeneous is bad. Administrative costs or other frictions in implementation might, of course, discourage the agency from setting the fully-differentiated vector of standards compatible with first-best. But giving a welfare-maximising EPA the power to tailor standards if it so wishes, the argument runs, cannot be welfare-reducing because *the worst the agency can do is opt not to use that power* (see, for two examples among many, Kolstad (1987) and Baumol and Oates (1988)). Parry and Williams (1999) provide a numerical evaluation of the potential loss due to uniformity.

We show - contrary to conventional wisdom - that banning tailoring is not necessarily bad. The model incorporates something reminiscent of, though distinct from, the ratchet effect widely recognized in the literature on strategic interaction between firms and regulatory agencies. The modern ratchet effect emerges from Weitzman (1980) - a regulated firm does not exert effort developing a new technology because it knows that the response of the regulator will be to ratchet up the requirements it faces. In that case the regulator might well want to commit not to engage in such ratcheting. The ratchet effect has been incorporated in models in all sorts of settings. Interesting recent applications include Meyer and Vickers (1997) and Dalen (1997, 1998). The wider literature on the benefits of 'tying the hands' of policymakers includes Panagariya and Rodrik (1993) and Haaporanta and Puhakka (1993). The related literature on under-investment in technological advance includes Laffont and Tirole (1996).

In the model that we present, when the regulatory agency has the opportunity to 'tailor' regulations, a subset of firms will not wish to show themselves to have access to a low cost compliance technology for fear they would be penalised in subsequent periods. When the regulated population is large, the requirement that regulation be uniform means that the stringency of the regulation faced by any *individual* firm will be independent of whether or not that firm has revealed its access to a low cost technology - if the regulator puts a requirement on one firm it must impose the same requirement on all others. Any *individual* firm with access to such a technology will, then, be willing to 'show its hand' by adopting in period 1.[2]

We establish that the level of welfare attainable by a welfare-maximising regulatory agency may be higher or lower under a regulatory constitution that requires uniformity as it would be under a constitution allowing the agency to treat different firms differently.

---

[2]We are by no means the first to point out that, in a dynamic principal-agent setting, (a) improving the information that a principal is allowed to use can adversely affect agents' incentives and (b) this effect may be dominant. Meyer and Vickers (1997) analyse the roles played by the ratchet effect and reputation effect in settings where there is scope for using indicators of *comparative* performance. Cremer (1995) shows that lowering the cost of aquisition of information about performance makes it more difficult for the principal to commit to threats and hence weaken incentives. The novelty here is in the benefits of a constitutional restriction that rules be uniform.

While a requirement that regulation be uniform does *not* 'tie the hands' of the agency - it doesn't allow or require it to commit to future standards - it *does* inhibit their freedom of movement. What is required from one firm must now be required from all the others.

The story relies, notice, on there being at least two - and perhaps a large number - of firms in the population. Untailored regulation then ties the stringency of regulation faced by any firm to characteristics of the population average. Because of this the model is distinct from those treating the regulation of a monopolist with unknown costs (Lewis (1996) synthesises this large literature). The effect identified is not, then, the standard ratchet effect that appears in those sorts of settings. Tying regulation to the population average characteristics can mitigate (in part) disincentives for adoption of low-cost technologies. The benefits of such mitigation must be weighed against the inefficiencies usually associated with uniformity.

The welfare results that we present will involve significant qualitative ambiguities, and it is worth flagging the conflicting welfare-impacts of tailoring at this early stage. From a welfare perspective outcomes under the two regimes - one in which tailoring is allowed, the other in which the agency is restricted to apply uniform standards - differ in two important ways. These are the static efficiency of requirements (the extent to which standards are set at the right level given the technology in use) and, second, the efficiency of the technology choices that are induced. More concretely we show that:
(a) When tailoring is not allowed all firms choose the least cost technology available to them in each period, which - other things being equal - is a good thing. When tailoring

is allowed, in contrast, technology choices will be efficient in period 2 (the last period) but period 1 incentives are distorted and technology choices are likely to be socially inefficient. (b) When tailoring is not allowed, *all* firms face inefficient standards in both periods - that is standards that do not equate the marginal cost of abatement to the firm with the marginal social benefit of abatement. In contrast, when tailoring is allowed some firms may face efficient standards in period 2.

The relative impact on welfare of these advantages and disadvantages (we will establish) is ambiguous, and depends upon cost conditions and other characteristics of the context. This points to an ambiguous normative appraisal of the requirement that regulatory requirements be uniform across firms.

If the power to tailor standards is not going to be used by the agency then it is ambiguous as to whether those powers should be taken away (this is *not* a case of redundant but irrelevant). The view that differentiated standards will dominate their uniform counterparts is not, in general, defensible. Even if the agency opts to use those powers, it may be welfare-improving to remove them.

It is worth contemplating in more detail two assumptions that we make regarding commitment.

Commitment to welfare maximization is assumed on the part of the regulator. This is a common assumption in models of this sort, even though it is understood that in reality changes in the incumbent government may lead to changes in the weights on emissions damage versus corporate costs that the agency applies. We can observe that various attempts have been made to institutionalize welfare-maximization - through requirements that agency rules satisfy a cost-benefit test, for example - in which case the analysis can be viewed as progressing on the assumption that such requirements are effective. Alternatively the assumption of continuing welfare maximization can be thought of as persistence in the average expected weights across time periods that involve multiple and a priori unpredictable political regimes. While other authors have analyzed how politicisation of regulatory agencies affects policy-making (e.g., McCubbins et al (1987), Glicksman and Schroeder (1991)) the focus of our paper is different. Our interest is in exploring how a welfare-motivated principal might 'manage' the game between regulators and firms in such a way as to ensure socially-desirable outcomes, abstracting from the possibility of 'taste' change on the part of the polity. This allows us to highlight the key strategic issues that arise in the game between regulatory agency and regulated parties, and maintain tractability.

Credible commitment to period 2 standards at the start of period 1 is assumed impossible, such that policy has to be sequentially-consistent. Again this is a common assumption in models of this sort. An alternative approach would be to assume that the regulator could write 'contracts' regarding the future environmental standards to which they would be held. The legal enforceability of such commitments is always likely to be questionable, however, with the government retaining the *de facto* flexibility either to dilute requirements, or to add to them, should the regulatory context change.[3]

---

[3]One of the ackowledged problems of the 'technology-forcing' approach to policy - setting very tough future targets in order to try to induce the regulated industry to innovate

We use a stylized two-period model to explore the role of the ratchet effect in a particular policy setting where the benefits of tailored regulation have been widely promoted. Contrary to conventional wisdom, barring the regulatory agency from tailoring standards may be good policy.

2. A Model

Consider a two-period model with no commitment. At the start of each period the regulatory agency (the 'EPA') sets an emissions standard for each of a large number of firms in a population. The emissions allowed by firm i will be denoted s_i, and exact compliance with the vector of standards is fully and costlessly enforced.[4]

Every firm has access to an established or 'back-stop' abatement technology which allows it to reduce its emissions to a level e at a cost

$$\theta_H.c(e)$$

- revolves around credibility of commitment. The regulated industry anticipates that if innovation fails to occur a forward-looking regulator will choose to reduce requirements, or postpone implementation deadlines, ex post. See Gerard and Lave (2005), for example, for insightful discussion in the context of the implementation of advanced emissions standards in the United States.

[4]A regulatory standard s_i is analogous to a non-tradeable pollution quota. We ignore other instruments. Note, also, that this is not a mechanism design problem, and that the EPA does not have the power to tax or spend, it simply sets standards.

where $c' < 0$, $c'' > 0$ (the marginal product of abatement spending is everywhere positive but diminishing). A fraction α of firms also have access to a better (lower cost) abatement technology which allows them to reduce emissions equivalently at a cost $\theta_L.c(e)$, $\theta_L < \theta_H$. 'Technology' can be interpreted broadly to include organisational reform, adjustment of work practices, product mix, relocation *etc.* We will refer to firms that do and do not have access to the alternative as L- and H-types respectively. Whilst H-type firms have no choice of technique, L-types can choose between the two technologies.

The regulatory agency is assumed unable to observe which firms in the population have access to the alternative technology (though it knows α). It can observe realised compliance costs, but those costs are not verifiable.[5]

[5]The assumption of non-verifiability is plausible wherever compliance is by means other than simple end-of-pipe solutions (*i.e.* in most settings). An inspector in an ongoing relationship with a firm and with 'insider' knowledge and experience of how that firm operates may be able to observe the amount of effort or distortion of practices (for example) that a firm exerts in ensuring compliance. He may well not readily be able to document that to a third-party. Because we are assuming that all the EPA does is set and enforce emissions standards the lack of verifiability is not particularly significant. If the EPA were empowered, in addition, to require firms to use the 'best available technology', and to then penalise retrospectively firms that didn't, then verifiability would matter. Combining such a power with an assumption of non-verifiability leaves the properties of

The agency minimises a social loss function

$$\sum_{t=1,2}\sum_{i}\left(d.s_i+\theta_i.c(s_i)\right), \tag{1}$$

the two-period sum of environmental damage and realised compliance costs. The analysis extends to unweighted sums of damage and cost, as can be seen by rescaling d. It is unable to commit to future standards (*i.e.* to period 2 standards at the start of period 1). Thoughout the paper we rule-out corner solutions by making the standard sort of assumption on abatement costs, more concretely here that $c'(0) > d/\theta_L$.

Our interest here is less to characterise how regulators will actually behave, more to say something about the type of regulatory 'constitution' under which they should be asked to operate. The objective here is to characterise optimal regulation under two scenarios - one in which the agency is free to set standards on a firm-by-firm basis, one in which it is obliged to set undifferentiated standards - and to compare the outcomes achieved. To reduce clutter in the paper we do not characterise optimal policy with commitment. It is straightforward to verify that with commitment and tailoring (that is where the EPA can commit to a pair of second period standards, assigned according to period 1 choice of technology) optimal policy involves distortion of standards in both periods.

---

the model unchanged.

2.1. Outcome when tailoring is prohibited

As a benchmark, consider the case in which tailoring is prohibited - the EPA is obliged to impose the same standard on every firm.

Every firm will use the least cost technology available to it (there is no incentive to try to hide type). In each period the EPA will set a uniform standard, s_U, to satisfy

$$-(\alpha.\theta_L + (1-\alpha).\theta_H).c'(s_U) = d, \tag{2}$$

i.e. to equate average marginal cost (the left hand side) with marginal damage (the right), with the former conditioned on efficient choice of technique. Two-period social loss under optimal uniform regulation can then be written

$$2.\left(d.s_U + (\alpha.\theta_L + (1-\alpha).\theta_H).c(s_U)\right) \tag{3}$$

which we will denote $L^{NO\,TAILORING}$. This provides a benchmark against which we compare performance when tailoring of standards is permitted.

2.2. Outcome when tailoring is allowed

Suppose, on the other hand, that the EPA is allowed to tailor standards.

We will start by characterising optimal standard-setting in period 2. If a firm was observed to use the low cost technology in period 1 then the EPA can infer that it is an L-type and will impose upon it the standard s_L that satisfies

$$-\theta_L.c'(s_L) = d. \tag{4}$$

If a firm was observed to use the high cost technology it may be *either* that it was an H-type firm that used the best (indeed only) technology available to it *or* that it was an L-type that chose to use the H technology for strategic reasons.

Let $\widehat{\pi}$ be the EPA's expectation of π, the proportion of L-types who used the L technology in period 1. Given that all firms will choose the least cost technique available to them in period 2, the optimal standard $s_{\pi}(\widehat{\pi})$ to impose upon a firm that was observed to use the H technology in period 1 is implicitly defined by:

$$-\left(f.\theta_L + (1-f).\theta_H\right) c'(s_{\pi}(\widehat{\pi})) = d, \tag{5}$$

where

$$f = \frac{\alpha.(1-\widehat{\pi})}{\alpha.(1-\widehat{\pi}) + (1-\alpha)} \tag{6}$$

is the probability that a firm observed to use the H technology was in fact an L-type. It is straightforward to verify that

$$\frac{\partial s_{\pi}(\widehat{\pi})}{\partial \widehat{\pi}} = -\left[\frac{(\partial f/\partial \widehat{\pi}).(\theta_L - \theta_H).c'(s_{\pi}(\widehat{\pi}))}{(f.\theta_L + (1-f).\theta_H).c''(s_{\pi}(\widehat{\pi}))}\right] > 0. \tag{7}$$

As $\widehat{\pi}$ goes up the EPA's expectation of the proportion of L-types using the H technology falls. This means that of the set of firms observed using the H technology a higher proportion really are now H-types, with no access to the cheaper alternative. The EPA therefore imposes a correspondingly less demanding standard.

Having characterised how the EPA will regulate in period 2, we can consider how a firm will make its choice of technique in period 1. An H-type has no option but to use the H technology. An L-type can use the L technology, revealing its type, and will face two-period costs

$$\theta_L.(c(s^1) + c(s_L)) \tag{8}$$

where s^1 is the first-period standard. Firms are *ex ante* identical and the EPA sets a uniform period 1 standard. In so doing it will take account of the incentives that that standard has for choice of technique. Note that this is not a procurement problem in which the EPA is allowed to 'buy' emissions reductions by setting a tax/transfer schedule. There are a variety of reasons - that we do not elaborate on here - why agencies may be precluded from the use of fiscal instruments of this kind. Alternatively it can use the H technology and face costs

$$\theta_H.c(s^1) + \theta_L.c(s_\pi(\widehat{\pi})). \tag{9}$$

Combining (8) and (9) tells us that a particular L-type will use the L technology - 'reveal' - if

$$c(s^1).(\theta_H - \theta_L) > \theta_L.(c(s_L) - c(s_\pi(\widehat{\pi}))). \tag{10}$$

In a rational expectations equilibrium (REE) $\widehat{\pi} = \pi$.

Proposition 1 *The REE exists, is unique and is stable.*

The proof of the three parts of Proposition 1 constitute Appendix 1. The proportion of L-types that reveal in this equilibrium will be denoted π^*, and for exposition we will let π' be the value of π^* in those situations in which a mixed equilibrium prevails. Since the two-period costs that result from L-types choosing the H versus L technology exactly balance in such an equilibrium, π' is implicitly defined by:

$$c(s^1).(\theta_H - \theta_L) = \theta_L.(c(s_L) - c(s_\pi(\pi'))) \tag{11}$$

where $s_\pi(\pi')$, the regulator's optimal choice of standard for firms using L technology in period 1, is in turn implicitly defined by

$$-\left(\frac{\alpha.(1-\pi')}{\alpha.(1-\pi')+(1-\alpha)}.\theta_L+\left(1-\frac{\alpha.(1-\pi')}{\alpha.(1-\pi')+(1-\alpha)}\right).\theta_H\right)c'(s_\pi(\pi'))=d. \quad (12)$$

Then

$$\begin{aligned}\pi^* &= \pi' \text{ if } 0<\pi'<1 \quad (13)\\ \pi^* &= 1 \text{ if } 1<\pi' \\ \pi^* &= 0 \text{ if } 0>\pi'\end{aligned}$$

As would be expected, π' is decreasing in s^1. In other words a less stringent emissions target in the first period leads to less L-types choosing to 'show their hand' by using the L technology in that period. π^*, then, is weakly decreasing in s^1. This is shown in Figure 1, in which the bold hatched line represents π^*. We can sensibly differentiate between three types of equilibrium. The EPA setting $s^1<\underline{s}$ generates a fully-separating equilibrium (all L-types adopt the L technology in period 1), setting $s^1>\bar{s}$ a pooling equilibrium (none adopt), and setting $\bar{s}>s^1>\underline{s}$ leads to partial separation and an equilibrium in which some L-types use their least cost technology in the first period while other L-types use the high cost technology.

First period social loss when tailoring is allowed can be written

$$d.s^1+\alpha.\pi^*(s^1).\theta_L.c(s^1)+(1-\alpha.\pi^*(s^1)).\theta_H.c(s^1), \quad (14)$$

The first term represents period 1 environmental damage. The second term is the compliance costs incurred by firms using the L technology. The third term is the compliance costs incurred by firms using the H technology.

Second period social loss can be written:

$$\alpha.\pi^*(s^1).(\theta_L.c(s_L) + d.s_L) + \tag{15}$$
$$(1 - \alpha.\pi^*(s^1)).d.s_\pi(\pi^*(s^1)) + \left[(1-\alpha).\theta_H + \alpha.(1 - \pi^*(s^1)).\theta_L\right].c(s_\pi(\pi^*(s^1))).$$

Here we have organised terms slightly differently. The first line is the total social cost (damage imposed plus compliance costs incurred) for firms known by the EPA to be L-types. The second line captures the total cost for the rest of the firms.

Ignoring discounting, the EPA chooses the first period standard to minimise two period loss. We will denote that optimal value s^{1*} and the implied minimised value of two period social loss $L^{TAILORING}$.

2.3. Comparing the Regimes - Allowing Tailoring Can Be Harmful

To reiterate, the outcome under the two regimes differs in two main ways: (a) When tailoring is not allowed all firms choose the least cost technology available to them in each period, which is a good thing. However, *all* firms face inefficient standards in both periods (standards that do not equate the marginal cost of abatement to the firm with the marginal social benefit of that abatement). In contrast, (b) when tailoring is allowed a subset of firms face efficient standards in period 2. Choice of technology is efficient in period 2 but will in almost all cases be *inefficient* in period 1.

The relative impact on welfare of these advantages and disadvantages (we will establish) is ambiguous, and depends upon cost conditions and other characteristics of the context. As the results we are looking to establish point to the existence of welfare ambiguities, it is sufficient to proceed to prove by examples. The parameterisation used is $c(s) = k - \ln(s)$ with $k = 1$, $\theta_L = 1$, $d = 4/5$, and all parameters subject to the constraint $c(s) > 0$ for all s used in any calculations. This parameterisation spans a slice in (α, θ_H)-space. Computing outcomes at each point in this slice through the parameter space yields results summarised in Figure 2. (The computational procedures involved in constructing the Figure are detailed in Appendix 2).

The parameter space is partitioned in two ways. Firstly, three types of equilibrium may result from the EPA's optimal strategy under tailoring. The solid curves in Figure 2 delimit these equilibrium types. Points in the region to the left of the central vertical curve correspond to combinations of parameters that imply $s^{1*} > \overline{s}$, so that under a regime allowing tailoring the agency will choose to implement a pooling equilibrium. To the right of that curve are combinations of parameters for which $s^{1*} < \overline{s}$, so that under a regime allowing tailoring equilibrium will involve separation. Full separation occurs in the upper part of this region, partial separation in the lower. Iso-π^* lines (associated with π^* equals 0.1, 0.3, 0.5, 0.7, and 0.9) are indicated by the dotted curves within the region of partial separation.

Secondly, social welfare when optimal policy is followed is sometimes highest when tailoring is permitted and sometimes highest when it is not permitted. Below the dashed curve, $L^{TAILORING} > L^{NO\,TAILORING}$, so it is best to prohibit tailoring. Above the dashed curve the inequality is reversed, so it is best to allow tailoring.

The example summarised in Figure 2 establishes two key results, for which intuition has already been provided.

Proposition 2 *When tailoring is permitted, optimal policy may involve the EPA setting the first-period standard such as to generate a pooling equilibrium, a partially-separating equilibrium or a fully-separating equilibrium. That is, s^{1*} may be greater than $\overline{s}$, less than $\underline{s}$, or between the two.*

This is proved by example by noting the existence of feasible parameter combinations that fall into each of the three regions bounded by the unbroken lines.

Proposition 3 *Social loss may be higher or lower when tailoring is permitted than when tailoring is prohibited.*

This is proved by example by noting the existence of feasible parameter combinations that fall above and below the broken line. Points on the line correspond to parameter combinations where the welfare outcomes are equated.

2.4. When Tailoring Is Harmful

Ideally one might determine whether allowing tailoring is harmful for specific contexts, and adjust the constraints put upon different agencies in different settings correspondingly. Some general principles can be drawn that help determine when allowing tailoring is likely to be good or bad.[6]

---

[6]These general principles presume the EPA is rational and acts to minimize social loss; in practice one must be sensitive to ways in which the regulator may be deviating from

It is always a bad sign if the regulator has the power to tailor standards but does not exercise that power. On the other hand, use of the power to tailor does not signal whether this power is bad or good. Formally, this is established by Propositions 4 and 5.

Proposition 4 *If the EPA has the power to tailor standards but does not use that power, then having that power is always welfare-reducing.*

This is proved as follows. The regulator never chooses $s_L = s_H$ since $c'(s_L) = -(d/\theta_L) \neq c'(s_H) = -(d/\theta_H)$. Hence the only case in which the EPA has the power to tailor standards but does not use that power is the case of complete pooling. With complete pooling the period 2 outcome is the same with or without the ability to tailor; the regulator chooses standard s_U. In the first period, however, the outcome with tailoring is unambiguously worse. With tailoring, L-type firms do not use the low-cost technology, and no choice by the regulator could do better than to achieve the first-period loss

$$\min_{s_1} \{d.s_1 + \theta_H.c(s_1)\}. \tag{16}$$

However, under a regulator that cannot tailor, some firms use the low-cost technology and the regulator achieves

$$\min_{s_1} \{d.s_1 + (\alpha.\theta_L + (1-\alpha).\theta_H).c(s_1)\}. \tag{17}$$

Since $(\alpha.\theta_L + (1-\alpha).\theta_H) < \theta_H$, the regulator that cannot tailor unambiguously achieves a lower social loss than the regulator that can tailor.

---

this practice since politics, regulatory capture, and bounded information and rationality might alter the ways tailoring is applied.

Proposition 4 says that this is not a case of 'redundant but harmless'. In many settings regulatory powers are not, in equilibrium, used, but still exert a positive influence on the characteristics of the outcome (*e.g.* the threat of a big penalty may prevent wrong-doing such that in equilibrium that penalty never has to be levied). Here the reverse applies. If the right to tailor is not being used, it should be removed.

Proposition 5 *If the EPA has the power to tailor standards and uses that power, then removing that power may increase or decrease welfare.*

Again, because we are establishing ambiguity an example suffices. Proposition 5 is proved by example by noting that the region in which optimal policy involves inducing separation is cut by the dashed curve.

More can be said about how the parameters affect the desirability of tailoring for the special case $c(s) = k - \ln(s)$. Numerical analyses over an extensive grid of parameter values indicates certain effects of parameter variation.[7]

Call allowing tailoring bad if $L^{TAILORING} > L^{NO\,TAILORING}$, and good if $L^{TAILORING} < L^{NO\,TAILORING}$. As k, d, and α increase, and as θ_L decreases (holding

---

[7]The grid of parameter values used consists of 11 equally-spaced values for each of the model's five parameters over the ranges $k \in [0.2, 2]$, $d \in [0.5, 2]$, $\alpha \in [0.01, 0.99]$, $\theta_L \in [0.2, 2]$, and $(\theta_H - \theta_L) \in [0.01, 2]$. Among the resulting $11^5 = 161,051$ sets of parameter values the limited set that implied negative costs - and so violated an assumption of the model - were excluded from comparisons. Computations were performed in the same way as for Figure 2.

$\theta_H - \theta_L$ constant), *ceteris paribus*, at critical thresholds allowing tailoring switches from bad to good but never *vice versa*. Thus, situations in which one would wish to allow tailoring occur when the fixed cost of environmental damage reduction is high, when the social cost of environmental damage is high, when a high proportion of firms have access to the low-cost technology, and when the total cost of environmental damage reduction is low.

This said, *within* regions in which the pooling result prevails $(L^{TAILORING} - L^{NO\,TAILORING})$ actually *rises* unambiguously as k, d, and α increase and as θ_L decreases (holding $\theta_H - \theta_L$ constant). Firms live under the fear of ratcheting regulation to such an extent that they pool completely. It is only by 'undoing' this fear sufficiently to move to a point where partial or full separation results under tailoring that $(L^{TAILORING} - L^{NO\,TAILORING})$ ceases to rise with these parameter changes, and hence that allowing tailoring becomes desirable.[8]

In delineating the harmful from non-harmful cases the essence of the trade-off is always between static and dynamic efficiency. In any single period tailored requirements are always (at least weakly) more efficient - for the usual set of reasons. The threat of tailoring in period two can, however, induce a subset of firms (those with low marginal costs) to make socially inefficient decisions in period one in order to hide their type. The terms of that trade-off depend upon the variety of parameters identified, and do

[8]Within regions in which $\pi > 0$, $(L^{TAILORING} - L^{NO\,TAILORING})$ is strictly non-increasing in k and d, and strictly increasing in θ_L (holding $\theta_H - \theta_L$ constant).

not allow for an unambiguous conclusion to be drawn regarding the benefits of allowing tailoring.

3. Conclusions

There is a large and growing body of literature which aims to compare the efficiency of alternative instruments of environmental regulation, and to characterise their optimal use. Weitzman (1974), Besanko (1987) and Stavins (1996) are well-known examples.

In this paper we have examined the comparative efficiency of two instruments - uniform emissions standards and firm-specific emissions standards - in a two-period model in which firms have private information about their access to low-cost abatement techniques. If the power to tailor standards is not going to be used by the agency then those powers should be taken away, as they are always harmful - this is not a case of redundant-but-harmless. Even if the agency opts to use those powers, it may be welfare-improving to remove them. In our application of environmental regulation, the results contrast with conventional wisdom amongst policy commentators that - absent issues of political economy, capture *etc.* - giving regulators more flexibility would always be good.

The analysis has been limited in scope, and contains a variety of simplifying assumptions. We have compared two particular instruments, namely uniform versus differentiated standards. We have not explored more sophisticated ways in which standards could be applied, such as offering firms a menu of period-1 standards contingent on the period-1 technique used, or using monetary transfers to induce revelation. Nor have we considered instruments other than standards.[9] What we refer to as 'optimal' policy in the body of the paper is always understood - as ever - subject to these self-imposed limitations.

We have also analyzed how the relative social loss, with and without tailoring allowed, changes with parameters. The analyses pertain to a particular functional form, involving costs of achieving an emissions standard s equal to $k - \ln(s)$. In contexts where at least some firms choose to exploit their access to better (low-cost) technology, the desirability of tailoring rises with the fixed cost of environmental damage reduction and the social cost of environmental damage, and falls with the total cost of environmental damage reduction. Tailoring also becomes desirable when the fraction of firms that have access to the low-cost technology is sufficiently high. The opposite trends are true

[9] There are political-economic or other considerations that might preclude the use of economic instruments. One potential line of extension would be to consider the utility of making period-1 standard conditional upon period-1 technique. In the current model we have assumed that whilst compliance costs are observeable to the EPA they are not necessarily verifiable. Because realised compliance costs (and hence technology choice) cannot be verified by a third-party, any sanction threatened against firms adhering to the 'wrong' standard would not, given typical legal institutions, be enforceable. If verifiability of realised compliance costs were assumed (*e.g.* if compliance is by means of 'off-the-shelf' solutions) the model would further be complicated by the need to consider the incentive for firms to 'pad' those costs.

in regions of the parameter space in which no firms exercise their ability to use better (low-cost) technology. This example suggests that care is needed in describing the sorts of settings in which discretion to tailor are likely to be beneficial versus detrimental.

The view that differentiated standards will dominate their uniform counterparts is not, in general, defensible on economic efficiency grounds. The environmental application is illustrative rather than exhaustive of the contexts in which the model might be applied.

4. Appendix 1: Existence, Uniqueness, and Stability of the REE

The three elements of Propositon 1 are established in three separate parts.

Stability of the equilibrium, with $\pi = \pi^*$, results from L-type firms' individual decisions to maximize financial benefits, given projected actions of the EPA. The EPA knows which L-type firms used L technology in period 1 (this is how it tailors), so that when setting period 2 standards it knows π, and firms make decisions in expectation of this knowledge. Firms' decisions to choose L or H production technologies in period 1 depend on the net benefit to choosing L over H.

The cost to an L-type firm will total $\theta_L \cdot c(s^1) + \theta_L \cdot c(s_L)$ if using the L technology in period 1, or $\theta_H \cdot c(s^1) + \theta_L \cdot c(s_\pi(\pi))$ if using the H technology in period 1. The net benefit to choosing L over H, which is a function of π, is therefore

$$B(\pi) = -\left(\theta_L \cdot c(s^1) + \theta_L \cdot c(s_L) - \theta_H \cdot c(s^1) - \theta_L \cdot c(s_\pi(\pi))\right). \qquad (18)$$

Its derivative is

$$B'(\pi) = \theta_L \cdot c'(s_\pi(\pi))s'_\pi(\pi) < 0. \tag{19}$$

An equilibrium $\pi^*(s^1)$ exists if it satisfies any of the following three cases: (i) $\pi = \pi^*(s^1) = 0$ and $B(\pi) \leq 0$, (ii) $\pi = \pi^*(s^1) = 1$ and $B(\pi) \geq 0$, and (iii) $\pi = \pi^*(s^1)$ with $0 < \pi^*(s^1) < 1$ and $B(\pi) = 0$.

Remark 1 *The equilibrium $\pi = \pi^*$ is stable, in that each L-type firm's incentive to choose L over H (and therefore increase π) has the same sign as $\pi^* - \pi$ if $\pi \neq \pi^*$ or $0 < \pi^* < 1$, or is nonpositive if $\pi = \pi^* = 0$ or nonnegative if $\pi = \pi^* = 1$.*

Proof: If $0 < \pi^* < 1$ then $B(\pi^*) = 0$ by (11). If $\pi^* = 0$ then $B(\pi^*) \leq 0$ and if $\pi^* = 1$ then $B(\pi^*) \geq 0$, using (18). Using (19) to analyze the value of $B(\pi)$ for π below or above π^*, $B(\pi)$ is positive when $\pi < \pi^*$ and negative when $\pi > \pi^*$. QED.

The following Lemma is used in proving uniqueness of the equilibrium. In the Lemma, (20) is (12) rearranged, using π rather than π' because the equation holds for any $\pi \in [0, 1]$ not just equilibrium values.

Lemma 6 *s_L and*

$$s_\pi(\pi) = c'^{-1}\left(-\frac{(1-\alpha\pi)d}{(1-\alpha)\theta_H + \alpha(1-\pi)\theta_L} \right) \tag{20}$$

are unique, and if s^1 is unique then $B(\pi)$ is unique.

Proof: Since $c'(s)$ is strictly increasing in s, its inverse $c'^{-1}(.)$ is also a strictly increasing function. Therefore solving for s_L in (4) yields a unique solution. The

EPA's second-period optimization problem to choose $s_\pi(\pi)$, in order to minimize the second line of (15), yields (20) as its first-order condition (the second-order condition is satisfied), and (20) is likewise unique for each π. By (18), $B(\pi)$ is unique given that s^1 is unique. QED.

Uniqueness and existence follows directly.

Remark 2 *Once the EPA chooses s^1, any equilibrium of firms' decisions, $\pi = \pi^*$, is unique.*

Proof: If $\alpha = 1$, the EPA sets standards uniformly as firms are all L-types, so the unique equilibrium is $\pi = \pi^* = 1$. The remaining case is $\alpha < 1$, and by Lemma 1, $B(\pi)$ is unique. Suppose an equilibrium of type (i) exists, and that also another equilibrium exists. For the other equilibrium $\pi > 0$ (since $\pi \in [0, 1]$), and since $B(0) \leq 0$ by (i) and $B'(\pi) < 0$ by (19), $B(\pi) < 0$ for all $\pi > 0$. This contradicts (ii) and (iii), so another equilibrium cannot exist. Similarly an equilibrium of type (ii) cannot coincide with another equilibrium, because $B(1) \geq 0$ by (ii) and $B'(\pi) < 0$ by (19), implying $B(\pi) > 0$ for all $\pi < 1$, contradicting (i) and (iii). Finally, an equilibrium of type (iii) cannot coincide with another equilibrium, because $B(\pi^*) = 0$ by (iii) and $B'(\pi) < 0$ by (19), implying $B(\pi) > 0$ for $\pi < \pi^*$ and $B(\pi) < 0$ for $\pi > \pi^*$, violating (i), (ii), and (iii) for any other equilibrium's value of $\pi \neq \pi^*$. QED.

Remark 3 *Once the EPA chooses s^1, an equilibrium of firms' decisions, $\pi = \pi^*$, exists.*

Proof: Suppose that no equilibrium $\pi = \pi^*$ exists. Then an equilibrium of type (i) does not exist, so $B(0) > 0$, and an equilibrium of type (ii) does not exist, so $B(1) < 0$.

Starting at $\pi = 0$, increasing π causes $B(\pi)$ to fall at the (finite) rate given by (19) until reaching $B(1)$. This implies that $B(\pi) = 0$ at some $\pi \in (0, 1)$, yielding an equilibrium of type (iii), a contradiction. Hence an equilibrium $\pi = \pi^*$ exists. QED.

The EPA's choice of s^1 is more complicated and we have not proven uniqueness nor existence. The social loss equals the sum of (14) and (15), which is a complicated function of exogenous parameters and of s^1. Uniqueness of s^1 will not hold if the social loss function happens to have two or more minima of equal depth. Existence of s^1 that minimizes the social loss would be assured if the social loss function were continuous and appropriate bounds set. The choice of s^1 can be bounded above and below by $s^1 \in [0, z/d]$ where $z = 2ds_L + 2\theta_H c(s_L)$ is an upper bound for the loss if $s^1 = s_L$, since choosing $s^1 > z/d$ yields a social loss strictly greater than z. However, the social loss function is discontinuous at

$$s^1 = c^{-1}\left(\left[c(s_L) - c\left(c'^{-1}(-d/\theta_L)\right)\right]\frac{\theta_L}{\theta_H - \theta_L}\right) \tag{21}$$

which is where π' switches sign causing π^* to switch between 0 and 1. Hence we have not ruled out the possibility that a well-defined optimal choice of s^1 might not exist in cases where the infinum of the loss occurs at this transition point, in which cases the EPA would be forced to choose an s^1 that yields a loss ε greater than the infinum for some small $\varepsilon > 0$. Such violations of existence and uniqueness can be treated as measure zero in the space of exogenous parameterizations. Since the EPA must choose a single s^1 regardless in advance of period 1, these issues have no material effect on outcomes of the model.

5. Appendix 2: Numerical Example

The results summarised in Figure 2 were calculated as follows. First, a random sample of 160,000 points was drawn using a joint uniform distribution of $\alpha \in [0, 1]$, $\theta_H \in [1.01, 2.5]$. At each point, firms' and the regulators' optimal decisions were determined by backward induction. Second, an iterative search procedure was used to solve for the value of θ_H at each boundary between regions for each value of α on a grid of 1001 evenly-spaced points from $\alpha = 0$ to $\alpha = 1$.

The backward inductive process is as follows. In period 2 all firms use the low-cost technology if possible. Let $\theta_U \equiv \alpha.\theta_L + (1 - \alpha).\theta_H$, $\theta_\pi \equiv f.\theta_L + (1 - f).\theta_H$ and $f \equiv \alpha.(1 - \pi)/(\alpha(1 - \pi) + (1 - \alpha))$. If not allowed to tailor the regulator sets a uniform standard $s_U = c'^{-1}(-d/\theta_U)$. If allowed to tailor the regulator sets standards of $s_L = c'^{-1}(-d/\theta_L)$ and $s_\pi = c'^{-1}(-d/\theta_\pi)$ respectively for firms that did and that did not use the low-cost technology in period 1.

In period 1, $\pi = 1$ if the regulator is not able to tailor. If the regulator is allowed to tailor, firms choose mixed strategies yielding $\pi^* = \max(0, \min(\pi', 1))$ where π' satisfies $c(s^1).(\theta_H - \theta_L) = \theta_L.(c(s_L) - c(s_\pi(\pi')))$. Using the functional form adopted, algebraic manipulation shows that $\pi' = 1 - \frac{f}{1-f}.\frac{1-\alpha}{\alpha}$ where $f = \frac{\theta_H - ds_\pi}{\theta_H - \theta_L}$, $s_\pi = \frac{s_L}{s_1^P} \exp(kP)$ and $P \equiv \frac{\theta_H - \theta_L}{\theta_L}$.

The regulator chooses in period 1 a standard for all firms. With no tailoring allowed $s^1 = s_U$. With tailoring allowed, the regulator chooses a value of s^1 that yields one of the three outcomes $\pi = 0$, $\pi = 1$ or $0 < \pi < 1$. For any hypothetical configuration of parameters, once the optimal choice is known for each of the three cases, the regulator's overall optimal choice is the one among the three that minimises two-period social loss. For the current functional form $\underline{s} = \exp(k)\left(\frac{\theta_L}{\theta_H}\right)^{\frac{1}{P}}$ and $\overline{s} = \exp(k)\left(\frac{\theta_L}{\theta_U}\right)^{\frac{1}{P}}$. Among choices that yield $\pi = 1$ the loss function is strictly convex in s^1 with a minimum at s_U so the regulator's optimal choice is $s^1 = s_U$ if $s_{U \leq} \underline{s}$ or $s^1 = \underline{s}$ otherwise. Among choices that yield $\pi = 0$ the loss function is strictly convex in s^1 with a minimum at s_H so the regulator's optimal choice is similarly $s^1 = s_H$ if $s_H \geq \overline{s}$ or $s^1 = \overline{s}$ otherwise. Among choices that yield $0 < \pi < 1$ the social loss function has not been proven to be everywhere convex, so an intensive grid search technique was used to find the lowest social loss associated with values of s^1 from $\underline{s}$ to $\overline{s}$. At each iteration of the grid search, 101 values of s^1 were tried, evenly spaced over the range being considered. Successive iterations searched within ± 1 grid point of the previously-found optimum until the optimal value of s^1 was determined to an acuracy of $\pm 10^{-6}(\overline{s} - \underline{s})$.

Loss functions were then compared with and without tailoring allowed, to determine the policy minimising two period social loss. This procedure was repeated for every point in the parameter space.

6. Bibliography

Ashworth, J. and I. V. Papps, 1991. "Equity in European Union Pollution Control", *Journal of Environmental Economics & Management* 20(1): 46-54.

Baumol, W. J. and W. E. Oates, 1988. *The Theory of Environmental Policy* (2nd edition), Cambridge University Press: Cambridge, England.

Besanko, D. A. 1987. "Performance vs Design Standards in the Regulation of Pollution", *Journal of Public Economics* 34(1): 19-44.

Boyer, M. and J.-J. Laffont, 1999. "Towards a Political Theory of the Emergence of Environmental Incentive Regulation", *RAND Journal of Economics* 30(1): 137-157.

Cremer, J., 1995. "Arms Length Relationships", *Quarterly Journal of Economics* 110(2): 275-90.

Dalen, D. M. 1998. "Yardstick Competition and Investment Incentives", *Journal of Economics & Management Strategy* 7(1): 105-26.

Dalen, D. M. 1997. "Regulation of Quality and the Ratchet Effect", *Journal of Regulatory Economics* 11(2): 139-155.

Freixas, X., R. Guesnerie and J. Tirole, 1985. "Planning Under Incomplete Information and the Ratchet Effect", *Review of Economic Studies* 52(2): 173-191.

Fuchs, E. P. and J. E. Anderson, 1987. "The Institutionalization of Cost-Benefit Analysis," *Public Productivity Review* 10(4): 25-33.

Gerard, D. and L. B. Lave, 2005. "Implementing Technology-Forcing Policies: The 1970 Clean Air Act Amendments and the Introduction of Advanced Emissions Controls in the United States", *Technological Forecasting & Social Change* 72(7): 761-779.

Glicksman, R. and C. H. Schroeder, 1991. "EPA and the Courts: Twenty Years of Law and Politics," *Law and Contemporary Problems* 54(4): 249-309.

Haaporanta, P. and M. Puhakka, 1993. "A Theory of Bureaucracy and Time-Consistency", *Economics & Politics* 5(3): 241-254.

Kolstad, C. D. 1987. "Uniformity vs Differentiation in Regulating Externalities", *Journal of Environmental Economics & Management* 14(4): 386-399.

Laffont, J.-J. and J. Tirole, 1996. "Pollution Permits and Environmental Innovation", *Journal of Public Economics* 62(1-2): 127-140.

Lichtenberg, E., R. Spear and D. Zilberman, 1993. "The Economics of Re-entry Regulation of Pesticides", *American Journal of Agricultural Economics* 75(4): 946-958.

McCubbins, M. D, R. Noll and B. R. Weingast, 1987. "Administrative Procedures as Instruments of Political Control," *Journal of Law, Economics and Organization* 3(2): 243-77.

Meyer, M. A. and J. Vickers, 1997. "Performance Comparisons and Dynamic Incentives", *Journal of Political Economy* 105(3): 547-581.

Panagariya, A. and D. Rodrik, 1993. "Political-Economic Arguments for Tariffs", *International Economic Review* 34(3): 685-703.

Parry, I. W. and Williams, W. H., 1999. "A 2nd Best Evaluation of Eight Policy Instruments in Reducing Carbon Emissions", *Resource & Energy Economics* 21(3d): 347-374.

Stavins, R. N. 1996. "Correlated Uncertainty and Policy Instrument Choice", *Journal of Environmental Economics & Management* 30(2): 218-229.

Weitzman, M. L. 1980. "The Ratchet Principle and Performance Incentives", *Bell Journal of Economics* 11(1): 302-308.

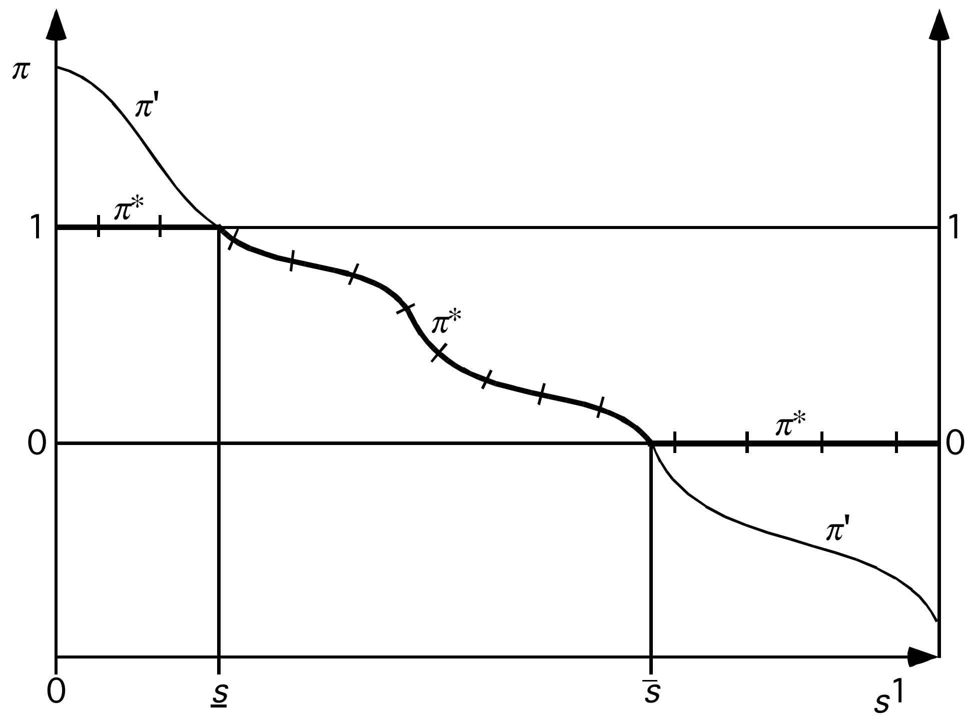

Figure 1. Equilibrium proportion of *L*-types using the low-cost technology in period 1, as a function of s^1.

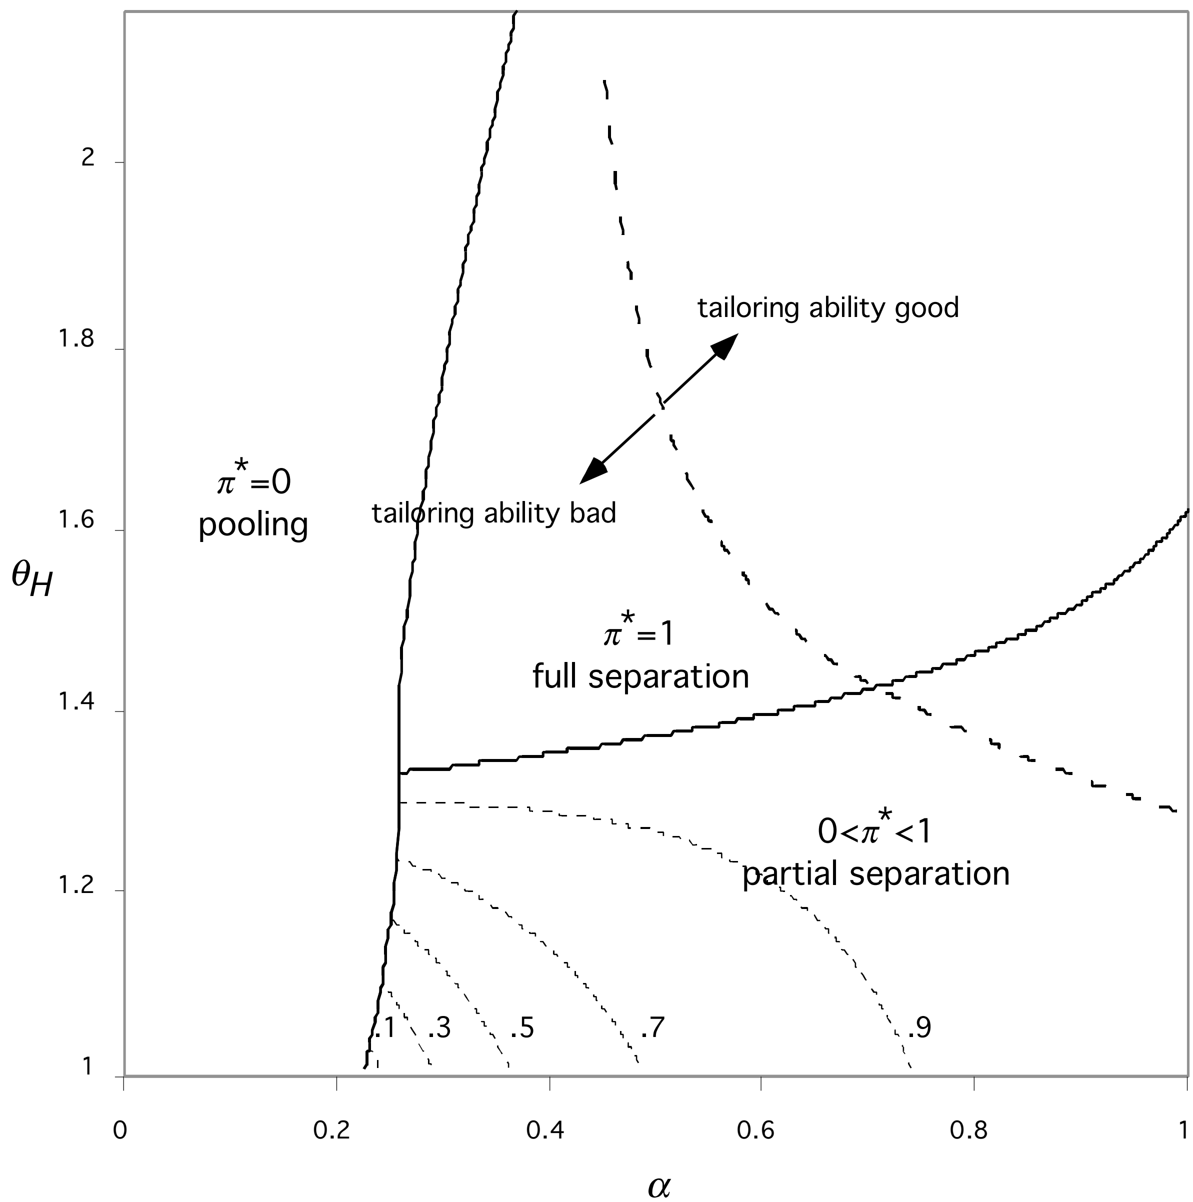

Figure 2. Outcomes of the model for a slice through the parameter space, using $c(s) = k - \ln(s)$, k=1, d=4/5, θ_L=1.