ECONSTOR Make Your Publications Visible.

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Hintermann, Beat; Lange, Andreas

Working Paper Learning abatement costs: On the dynamics of optimal regulation of experience goods

CESifo Working Paper, No. 4058

Provided in Cooperation with: Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Hintermann, Beat; Lange, Andreas (2012) : Learning abatement costs: On the dynamics of optimal regulation of experience goods, CESifo Working Paper, No. 4058, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at: https://hdl.handle.net/10419/69570

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



Learning Abatement Costs: On the Dynamics of Optimal Regulation of Experience Goods

Beat Hintermann Andreas Lange

CESIFO WORKING PAPER NO. 4058 CATEGORY 10: ENERGY AND CLIMATE ECONOMICS DECEMBER 2012

An electronic version of the paper may be downloaded • from the SSRN website: www.SSRN.com • from the RePEc website: www.RePEc.org • from the CESifo website: www.CESifo-group.org/wp

Learning Abatement Costs: On the Dynamics of Optimal Regulation of Experience Goods

Abstract

We study the introduction of new technologies when their costs are subject to idiosyncratic uncertainty and can only be fully learned through individual experience. We set up a dynamic model of clean experience goods that replace old polluting consumption options and show how optimal regulation evolves over time. In our base setting where social and private learning incentives coincide, the optimal tax of the polluting consumption is increasing over time. However, if social and private learning incentives diverge, we show that it will be optimal to temporarily increase the tax rate beyond net marginal external damages to induce optimal learning, before reducing the tax rate to the steady-state level. Alternatively, one needs to complement the tax by subsidies for first-time users which will be phased out over time. Similar results apply if consumers have biased expectations. We therefore give a rationale for introductory subsidies of new, clean technologies and non-monotonic tax paths from a perspective of consumer learning.

JEL-Code: D820, L510, H210, H230, Q520, Q550.

Keywords: experience goods, dynamic regulation, learning by doing, new technology, externalities, pollution.

Beat Hintermann University of Basel Department of Public Economics and Policy Peter Merian-Weg 6 Switzerland - 4002 Basel b.hintermann@unibas.ch

Andreas Lange University of Hamburg Department of Economics Von-Melle-Park 5 Germany - 20146 Hamburg Andreas.Lange@wiso.uni-hamburg.de

1. Introduction

Many environmental regulation schemes involve taxes or subsidies that change over time. In this paper, we study the dynamics of environmental regulation to control the adoption of a socially beneficial experience good. That is, by trying out a new, less polluting consumption choice, consumers may learn about their specific individual net costs of its use.

Examples are widespread: car users are often only partially informed about the specific costs and benefits of using public transport (or other means of transportation). Ecologically produced food and clothing may have attributes unknown to the inexperienced user, including taste, durability, social acceptance and the like. Household or office appliances are often advertized in terms of their higher energy efficiency relative to the older versions they seek to replace, but consumers have only limited knowledge about operating costs and convenience associated with the new products. Experience learning may also relate to the sensation of warm-glow or social reputation from using environmentally friendly products (Andreoni 1990, Benabou and Tirole 2006). Common to these examples is that the new product reduces an externality and that consumers are uncertain about their personal benefits and costs before trying out the new product.

We study the optimal regulation of an experience good and explore rationales for initial subsidies for a new product, and demonstrate how optimal regulation levels change over time. We focus on the dynamics of government intervention that is driven exclusively by consumers' personal learning from adopting a new technology. That is, we abstract both from learning spillovers among consumers as well as from supply-side arguments like cost reductions through an intensified use of the new technology that may be caused by technological spillovers among firms.

The study of experience goods is novel in the environmental economics domain. Our paper is, however, related to studies in industrial organization that focus on monopolistic supply of experience goods. Bergemann and Välimäki (2006) examine monopolistic pricing of experience goods in a dynamic model. They show that the dynamics crucially depend on a simple dichotomy between mass and niche markets.¹ While prices in mass markets are declining over time, they may

¹ In their paper, a mass market is characterized by the optimal price for informed consumers being low enough such that uninformed consumers have an incentive to try the new product. In such a market, all consumers will eventually

initially be low but subsequently increase in niche markets. The low initial prices are set with a focus on increasing market penetration, whereas the higher prices in the steady state (where no more learning takes place) maximize monopoly profits. The literature on dynamic pricing of experience goods goes back to Shapiro (1983) who considers learning in a simple two-period model. Pricing of experience goods has also been discussed by Cremer (1984), Farrell (1986), Milgrom and Roberts (1986), and Tirole (1988).

We set up a dynamic model of a new experience good that replaces an old consumption option generating an environmental externality. We consider an infinite-horizon, discrete-time model with a continuum of consumers, who have (at most) unit demand per period for the new product. Alternatively, they consume the traditional, more polluting product. Consumers differ in their expected net costs of using the new technology and face an ex ante unknown cost component. The latter is subject to idiosyncratic uncertainty and can only be learned through individual experience. We assume that consumers learn their true costs/valuation when first consuming the new technology for one time period.²

We consider two different regulatory regimes: we start by analyzing the first-best case in which the regulator, at each point in time, determines both the number of inexperienced consumers that are exposed to the new technology for the first time, and the set of experienced consumers who should continue using the technology. Second, we consider a setting where the regulator needs to rely on subsidies/taxes only. Here, the subsidy in the given period determines both the behavior of the experienced consumers as well as the inexperienced consumers' decision to try the new technology.

We show how optimal regulation evolves over time. As long as consumers' and social discount rates coincide and expectations about net costs of the new technology are unbiased, the first-best case can be decentralized by relying on a tax for the polluting technology only, with the

learn. Conversely, a niche market will involve consumers that do not try the new product as they do not have an incentive to do so at the optimal price for informed consumers.

² This clearly is a highly stylistic assumption. In general, learning takes places at different rates and consumers may have different initial knowledge about the new technology. Empirical evidence suggests that the rate of learning crucially depends on the market (e.g., Erdem and Keane (1996), Ackerberg (2003), Erdem, Imai, and Keane (2003), Israel (2005), Osborne (2005), and Goettler and Clay (2006), Crawford and Shum (2005)).

optimal tax rate increasing over time. If private discount rates exceed the social discount rate or consumers cost expectation is biased, however, the first-best solution can only be implemented by complementing the tax with a subsidy for first-time users. This subsidy will be decreasing over time. If the regulator cannot discriminate between first-time and experienced users, the second best taxation scheme may involve a non-monotonic path: tax rates are first increasing with a focus on reaching the optimal amount of learning, before being reduced to the level that reflects the marginal social costs of using the polluting alternative.

Our results are qualitatively similar to those derived by Bergemann and Välimäki (2006) in the context of optimal monopoly pricing of experience goods. During the approach path, the government (or monopolist) sets the tax with a focus on inducing optimal learning. In the steady state, however, no more learning takes place, and the tax (price) is chosen in order to maximize social welfare (monopoly profits).

We believe that our results indicate an important reason for a slow introduction of pollution taxes that is motivated by the fact that consumers are uncertain about their personal costs or benefits from using a new and cleaner technology. Only by trying it will they learn about the personal fit. We show that such a setting not only motivates taxes that are increasing over time, but also that it may require introductory subsidies for first-time users that are phased out over time.

Our paper is structured as follows. Section 2 presents our basic model and section 3 discusses the social optimum. Section 4 considers the case of first-best regulation, while we turn to a second best regulation that solely relies on usage taxes in section 5. Section 6 concludes.

2. Base model

Consumers choose one of two mutually exclusive alternatives, which we will label A and B. Alternative A is the status quo and all of its costs and benefits are known, including environmental externalities (example: commuting by car). Alternative B is characterized entirely by private costs (example: commuting by train). For simplicity, we assume that alternative B does not have external costs or benefits, and that alternative A does not generate private costs.³ Consumers who have never used alternative B are uncertain about the total private costs and benefits involved, which varies by individual. Once a consumer tries B for one period, she learns her true net costs (or equivalently, her overall utility). If these net costs are below those associated with alternative A, she will continue using alternative B; otherwise she will revert to alternative A.⁴ Consumers ignore external effects when maximizing their utility.

We separate the private net costs into a part that is unknown to the consumer until she tries alternative B, and another part that she knows beforehand. This allows for the possibility that consumers have some information about some costs, even if they are unsure about the total costs.

We will denote a consumer's ex-ante known costs as Δ , which are distributed within the population according to the p.d.f. $g(\Delta)$ and c.d.f $G(\Delta)$. In addition, each consumer faces a cost component δ that is unknown before experiencing the good. We denote the p.d.f. and c.d.f. of δ across the population as $f(\delta)$ and $F(\delta)$, respectively. We assume that Δ and δ are distributed independently of each other and have full support. Without loss of generality, we normalize $E[\delta] = 0$ such that a consumer's type can be characterized by her ex ante expected costs $E(\theta | \Delta) = E(\Delta + \delta) = \Delta$.

Every period, a fraction of consumers tries alternative B for the first time. The distribution of Δ defines the sequence according to which consumers try alternative B, with the first ones to learn being consumers with the lowest expected costs Δ . This means that we can express the fraction of consumers who have learned at any given time as $G(\Delta_t)$, where Δ_t refers to the threshold in expected private costs below which all consumers have learned.

³ This assumption conveys no loss of generality as long as the marginal costs of B are larger than those of A even for the first user that switches from A to B. In this case, defining C=C(B)-C(A) yields the exact same first-order conditions as in our model. We can therefore interpret C as the incremental costs from providing alternative B relative to A, which are assumed to be positive, increasing, and convex. The same logic applies to environmental damages.

⁴ Key to our model is the reversibility of the decision. For our example of car commuting vs. public transportation, this means that consumers would keep the car when trying the public transport option rather than selling it immediately. Other examples could include renting of zero-energy housing, or of electric cars. Again, the key assumption is that the decision is reversible without additional costs, that the new technology is costlier than the old one, and that switching from the old to the new technology reduces an externality.

We denote θ_t as the threshold for private net costs up to which *informed* consumers continue to use alternative B at time t. A consumer who has already learned her private net costs $\theta = \Delta + \delta$ by trying alternative B chooses to use B if and only if $\theta = \Delta + \delta \le \theta_t$. Let $\Omega_t \in [0,1]$ describe the fraction of consumers that use B at time t. Normalizing the mass of consumers to 1, we obtain

$$\Omega_t \equiv \int_{-\infty}^{\Delta_t} F(\theta_t - \Delta) dG(\Delta) + G(\Delta_{t+1}) - G(\Delta_t)$$
(1)

where the first term describes the fraction of informed consumers that chooses B over A at any given time *t*, and $G(\Delta_{t+1}) - G(\Delta_t)$ is the fraction of the consumers that learn in period t.⁵

The use of the dirty alternative A generates external damages $D = D(1-\Omega_t)$ (with D' > 0, $D" \ge 0$). The clean alternative B causes operating costs that depend on the number of users, $C = C(\Omega_t)$ (with C' > 0, $C" \ge 0$). We assume that in any period, the usage rate of alternative B is capacity constrained: $\Omega_t \le K_t + k_t$ where K_t denotes the already existing capacity from the previous period and k_t its expansion in period *t*. Expanding the capacity of alternative B in period *t* to $K_{t+1} = K_t + k_t \ge K_t$ generates costs $A(k_t)$ (with $A' \ge 0$, A'(0) = 0, A" > 0). A natural example for our study is that of building infrastructure, e.g. for a public transportation network. Such networks are costly to build, and it is usually a slow process such that costs per period are highly convex.

3. The social optimum

We set up a dynamic optimization problem from the point of view of the social planner who aims to minimize costs (which is equivalent to maximizing social welfare under a utilitarian social welfare function). Defining the per-period increment in the fraction of informed consumers as $s_t = \Delta_{t+1} - \Delta_t$ and substituting into the objective function, the intertemporal optimization problem

⁵ Because the threshold θ_t can change over time, a consumer could potentially switch back and forth between the two alternatives.

can be written as minimizing social costs (SC) by choosing the control variables θ_t , s_t and k_t at every point in time, subject to state equations for the state of learning and the available capacity, a capacity constraint, nonnegativity constraints and initial conditions for the state of learning and capacity:

$$\begin{split} \min_{\theta_{t},s_{t},k_{t}} \quad SC &= \sum_{t=0}^{\infty} e^{-rt} \cdot \left(\int_{-\infty}^{\Delta_{t}} \int_{-\infty}^{\theta_{t}-\Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_{t}}^{\Delta_{t+1}} \Delta dG(\Delta) + D(1 - \Omega_{t}) + C(\Omega_{t}) + A(k_{t}) \right) \\ s.t. \quad \Delta_{t+1} &= \Delta_{t} + s_{t} \\ K_{t+1} &= K_{t} + k_{t}; \\ K_{t+1} &= K_{t} + k_{t}; \\ \theta_{t} \geq 0; s_{t} \geq 0; k_{t} \geq 0; \\ \Delta_{t=0} &= \Delta_{0}; K_{t=0} = K_{0} \end{split}$$
(2)

The first term measures the net costs of informed consumers that choose alternative B at period t; the second term reflects the expected per-period costs of learning; and the last three terms refer to net environmental damages, net costs of the clean alternative and the costs of capacity extensions, respectively. For each period, the social planner determines the cost cutoff value θ_t for informed consumers as well as rate of first-time users s_t and the capacity extension k_t .⁶ The Bellman equation corresponding to (2) is

$$J\left(\Delta_{t}, K_{t}\right) = \min_{s_{t}, \theta_{t}, k_{t}} \quad SC + e^{-r} \cdot J\left(\Delta_{t} + s_{t}, K_{t} + k_{t}\right)$$
(3)

subject to the capacity and nonnegativity constraints and initial conditions from (2), leading to a Lagrangian of the form

$$L = SC + e^{-r} \cdot J\left(\Delta_t + s_t, K_t + k_t\right) + \lambda_t (K_t + k_t - \Omega_t)$$

$$\tag{4}$$

⁶ Per-period capacity extension k_t , and thus also the state of extension K_t , depend on the choice of θ_t and s_t over time and could be substituted out, reducing the problem to one with two control variables. Specifically, $k_t = \max(\Omega_t - \Omega_{t-1}, \theta)$. We chose to leave it in the model because it facilitates exposition.

Dropping the function arguments for $D(\cdot), C(\cdot)$ and $A(\cdot)$, the optimality conditions are

$$(D'-C'-\lambda_t)-\theta_t \ge 0; \qquad \qquad \theta_t \ge 0$$
(5)

$$\left[\Delta_{t+1} - (D' - C' - \lambda_t)\right] \cdot g(\Delta_{t+1}) + e^{-r} \frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial \Delta} \ge 0; \qquad s_t \ge 0$$
(6)

$$A_{t}^{'} - \lambda_{t} + e^{-r} \frac{J(\Delta_{t+1}, K_{t+1})}{\partial K} \ge 0; \qquad k_{t} \ge 0$$
(7)

$$K_t + k_t - \Omega_t \ge 0; \qquad \qquad \lambda_t \ge 0 \tag{8}$$

with complementary slackness holding everywhere. A solution to the Bellman equation is a sufficient condition for optimality. Equation (5) implies that along the optimum approach path as well as in the steady state, marginal private cost θ_i associated with alternative B must be equal to marginal net benefits $(D'-C'-\lambda_i)$ of using this alternative, where λ_i adjusts for the current-period value of a marginal capacity expansion. Equation (6) equates the net social cost of subjecting the marginal uninformed consumer (the consumer with $\Delta = \Delta_{i+1}$) to learning in period *t*, consisting of the marginal expected cost of using this alternative Δ_{i+1} less the marginal reduction in damages, to the discounted marginal social value of having an additional informed consumer in the next period. If learning is not constrained by current capacity such that $K_i > \Omega_i$, marginal capacity expansion has no value, in which case (8) implies that $\lambda_i = 0$. For $\lambda_i > 0$, equation (7) states that marginal expansion costs A'_i have to be equal to the current-period marginal value of expansion λ_i plus the discounted social marginal value of having a larger capacity in the next period.

The intertemporal arbitrage conditions are given by

$$J_{\Delta}(\Delta_{t}, K_{t}) = \int_{-\infty}^{\theta_{t} - \Delta_{t}} (\Delta_{t} + \delta) dF(\delta) g(\Delta_{t}) + \Delta_{t+1} g(\Delta_{t+1}) - \Delta_{t} g(\Delta_{t})$$

$$- (D' - C' - \lambda_{t}) \Big[F(\theta_{t} - \Delta_{t}) g(\Delta_{t}) + g(\Delta_{t+1}) - g(\Delta_{t}) \Big] + e^{-r} J_{\Delta}(\Delta_{t+1}, K_{t+1})$$

$$J_{K}(\Delta_{t}, K_{t}) = -\lambda_{t} + e^{-r} J_{K}(\Delta_{t+1}, K_{t+1})$$

$$(10)$$

We can use these equations to explore the dynamics of the optimal approach path by means of a phase diagram that relates the fraction of informed consumers Δ_t to the cost cutoff θ_t in any given period. Because an interior first-best solution will feature a monotonic approach path for all variables, conditions (5)-(8) will hold with equality during the approach path as well as in the steady state. This property can be used to combine eqs. (6)-(10) (see Appendix) to get

$$e^{r}\left(\theta_{t}-\theta_{t-1}\right) = \left(e^{r}-1\right)\left(\theta_{t}-\Delta_{t}\right) + \int_{-\infty}^{\theta_{t}-\Delta_{t}}\left(\theta_{t}-\Delta_{t}-\delta\right)dF(\delta)$$

$$(11)$$

$$\lambda_{t} = A_{t} - e^{-r} A_{t+1}$$
(12)

A steady state is given by (5), (11) and (12) with $s_t = k_t = 0$ and $\theta_t = \theta_{t-1}$.⁷ To analyze the solution in Δ_t / θ_t -space, we derive the equation of motion for Δ_t by setting $s_t = 0$ in eq. (5)⁸ while holding λ_t constant,⁹ and the equation of motion for θ_t by setting $\theta_t = \theta_{t-1}$ in eq. (11).

Starting with the former we obtain an implicit relationship between θ_t and Δ_t which we denominate $\tilde{\theta}(\Delta_t)$, and which can be shown to be decreasing (see Appendix):

⁷ For $A'_{l0}=0$, eq. (11) implies that $\lambda_t=0$ for $k_t=k_{t+1}=0$. But with no learning and no change in the cost threshold, the same number of users will continue to use alternative B in the following period such that no capacity expansion is necessary. With a monotonic approach path, it follows that $k_{t+1}=0$ if $s_t=k_t=0$ and $\theta_t=\theta_{t+1}$.

⁸ The learning rate s_t appears in the usage rate Ω_t as defined in eq. (1), which consists of informed users and new users ("learners"). Setting $s_t=0$ implies that usage of the clean alternative is restricted to informed users.

⁹ A completely rigorous treatment would leave λ_t free and result in a 3-dimensional phase diagram. However, we know that $\lambda_t=0$ at the steady state such that the third dimension would complicate the exposition with little gain in intuition.

$$\left. \theta_t(\Delta_t) \right|_{s_t=0;\lambda_t=\bar{\lambda}} \equiv \tilde{\theta}(\Delta_t); \qquad \frac{d\tilde{\theta}(\Delta_t)}{d\Delta_t} = \frac{-(D"+C")F(\theta_t - \Delta_t)g(\Delta_t)}{1 + (D"+C")\int_{-\infty}^{\Delta_t} f(\theta_t - \Delta)dG(\Delta)} < 0$$
(13)

The economic interpretation of (13) is as follows: When no more learning takes place, the presence of more informed agents leads to higher usage rate of the clean technology, all else equal, and thus to lower marginal social costs. But this means that the cutoff value θ_t that fully internalizes external costs (for any given value of λ_t) will be lower. Conversely, if fewer people have learnt, usage of the clean technology will decrease, and a higher θ_t is needed to ensure that (5) holds.

Setting $\theta_t = \theta_{t-1}$ in equation (11) and solving for θ_t leads to another relationship between cutoff costs and the state of learning, which we denote by $\hat{\theta}(\Delta_t)$. It can be interpreted as the threshold of social marginal costs θ_t of using the old technology at which consumers up to ex ante expected costs of Δ_t should optimally learn. Intuitively, we obtain an increasing relationship between Δ_t and θ_t (see Appendix):

$$\left. \theta_t(\Delta_t) \right|_{\theta_t = \theta_{t-1}} \equiv \hat{\theta}(\Delta_t); \qquad \frac{d\hat{\theta}(\Delta_t)}{d\Delta_t} = 1$$
(14)

Hereby, the difference $\Delta_t - \theta_t$ can be interpreted as the value of learning for a consumer with ex ante expected costs Δ_t when regulation stays at θ_t . It is constant in Δ_t due to our specification that a consumer's ex post costs for are given by $\theta = \Delta + \delta$.

The two relationships are shown graphically in Figure 1. The steady state is defined by the intersection of the lines defined by $\hat{\theta}(\Delta)$ and $\tilde{\theta}(\Delta)$, where no learning takes place and θ_t remains constant. Substituting (1) into (5) implies that for $\theta_t < \tilde{\theta}(\Delta_t)$, learning will be positive ($s_t > 0$)

such that the system moves to the right below $\tilde{\theta}(\Delta_t)$, and to the left above.¹⁰ From (11), we see that $\theta_t > \theta_{t-1}$ if $\theta_t > \hat{\theta}(\Delta_t)$ and vice versa. These dynamics are illustrated by the arrows of motion in the figure. As a consequence, the optimal path into the steady state is characterized by cut off costs θ_t that increase over time and, naturally, by increasing Δ_t with consumers with the smallest expected costs trying the new technology first. The capacity of the systems K_t increases simultaneously. The dashed line in Figure 1 illustrates an example for the path into the steady state.

For the first-best solution, we can therefore summarize this result as follows:

Proposition 1

In the first-best transition path into the steady state, the marginal cost of participation $\theta_t = D'(1 - \Omega_t) - C'(\Omega_t) - \lambda_t$ increases over time.

This result may be surprising at first glance: One might expect that as consumers learn, more people use the new technology such that net marginal damages D' - C' decrease over time (recall that D' - C' is declining in the usage rate Ω_t). Proposition 1 implies that this is only the case if the marginal value of capacity expansion declines by even more. With sufficiently low expansion costs, $(-\lambda_t)$ will increase by less than θ_t , such that D' - C' will increase too, implying a *declining* usage rate. This rather counterintuitive result for a new technology can be explained as follows: If capacity expansion is free or very cheap such that the shadow value of capacity is at or near zero in the beginning (recall that it is always zero in the steady state), it will be optimal for consumers to learn early on, because this enables them to take advantage of the new technology should their idiosyncratic cost $\Delta + \delta$ turn out to be low enough ($\theta_t \ge \delta + \Delta$). Increasing net environmental costs D' - C' therefore imply that the increase in informed users due to an increase in θ_t is over-

¹⁰ This latter point applies in theory, but since we do not allow for "unlearning" in our model, a disequilibrium above the $\tilde{\theta}(\Delta_t)$ -line would have to be adjusted by lowering θ_t instead. Note also that due to the nonnegativity condition on learning, the initial state of learning Δ_0 must be below the steady state in order for this solution to hold. For a higher initial state of learning, no additional learning should take place and we obtain a trivial dynamics of the system. Similarly, the initial capacity K_0 has to be below steady-state capacity in order for the capacity constraint to be binding.

compensated by the decline in first-time users.¹¹ Note also that a declining usage rate implies building too much capacity, which cannot be optimal if expansion is costly.

Next, we examine how the steady state depends on the underlying parameters of the model. The relationship $\hat{\theta}(\Delta_t)$ depends on the distribution of $f(\delta)$ (and on whether consumers form unbiased expectations of personal costs $\theta = \Delta + \delta$, a point to which we return below). It further increases in r, which is a measure of consumers' (im)patience: Higher discounting implies a larger marginal cost threshold to make learning worthwhile for the same agent, all else equal. With an increase in r, the steady state moves up the $\tilde{\theta}(\Delta_t)$ -line leading to a larger θ^{SS} , combined with a smaller steady-state Δ^{SS} . This effect is shown in Figure 2 as a move from $\hat{\theta}_0(\Delta_t)$ to $\hat{\theta}_1(\Delta_t)$, leading to a new steady state SS'.

The function $\tilde{\theta}(\Delta_t)$ increases with marginal social damages D'-C'. A more polluting original technology (or equivalently, a cleaner new technology) increases the social value of learning such that $\tilde{\theta}_0(\Delta_t)$ is shifted to $\tilde{\theta}_1(\Delta_t)$, the steady state moves up along the $\hat{\theta}_0(\Delta_t)$ line to *SS''*, which is associated with both a higher state of learning as well as a higher θ^{SS} .

We illustrate our result numerically, using the following functional forms: The uniform distribution with support [-0.5, 0.5] for $f(\cdot)$; the uniform distribution with support [0.5, 1.5] for $g(\cdot)$;¹² $D(1-\Omega_t) = (\alpha/2) \cdot (1-\Omega_t)^2$ and $C(\Omega_t) = (\beta/2) \cdot \Omega_t^2$, and $A(k_t) = (\gamma/2) \cdot k_t^2$.

¹¹ This follows from the definition of Ω_t : If D'-C' increases, Ω_t has to decrease. But with increasing θ_t and Δ_t , usage by informed consumers (the first part in eq. (1)) has to increase, such that the reason for the decline in the usage rate has to be a decrease in new users (the second part in (1)).

¹² The rationale for these choices is primarily tractability. With the uniform distribution of unit support, $f(\delta)=g(\Delta)=1$, such that $dF(\delta)=d\delta$ and $dG(\Delta)=d\Delta$. The choice of $f(\delta)$ further ensures that $E[\delta]=0$. We set the lower limit of $g(\Delta)$ to 0.5 in order to avoid negative tax rates. The fraction of informed consumers at any moment is given by $G(\Delta_t)=\Delta_t$ -0.5.

Figure 3 shows the time paths of θ_t , Δ_t , s_t , K_t , k_t and Ω_t for various choices of γ , with α , β and *r* held constant (the qualitative nature of the paths is stable across these parameter values). The marginal cost θ_t , the state of learning $\Delta_t - 0.5$ and total capacity K_t increase, whereas the rate of capacity expansion k_t decreases over time. The participation rate Ω_t decreases if capacity expansion is cheap (γ at or close to zero), because the decrease in initially very high learning rates overcompensates the increase in informed users. The higher the expansion costs, the slower are the increases of θ_t , Δ_t and Ω_t . The rate of learning rates s_t monotonically decreases with low expansion costs, but is inversely U-shaped with medium to high costs.¹³

4. Optimal policy choices

We now address the question of how this optimal transition path can be decentralized by price instruments. For this, we consider a tax (price) τ_t on using the dirty alternative in period t as well as a subsidy¹⁴ σ_t for first-time users of the clean alternative, i.e. those consumers who learn about their personal cost δ_i in period t.

Consumers choose if and when to learn (i.e. to try the clean alternative for the first time), based on their private known costs Δ_i , their expected unknown costs, and the prevailing tax and subsidy rates. Conditional on the policy-path (τ_t , σ_t) and their private rate of discount r_p , consumers choose the optimal moment of learning $T(\Delta)$ that minimizes the current value of expected total costs by solving the stopping problem

¹³ The intuition behind the nonlinear shape is as follows: Suppose that expansion is costly such that only a small fraction of consumers should learn in the first period. Due to the low cutoff value, most of these first-period learners will revert to the old technology, leaving room for more learning using the capacity built in the first period. Together with the capacity added in period 2, learning will be higher in the second period than in the first. As θ_t increases, the fraction of experienced to first-time users will also increase, slowing the increase in learning and eventually leading to a falling s_t .

¹⁴ Theoretically, the subsidy could be negative, i.e. a tax. It corrects for a learning inefficiency: If consumers are too impatient (meaning that rP>r), they require a learning subsidy to learn according to (11). In contrast, if they are too eager to learn, they have to be deterred from too fast learning with a learning tax. We refer to a subsidy because it is difficult to motivate the situation of rP<r, but the model does not exclude the possibility that $\sigma t<0$.

$$\min_{T} \sum_{t=0}^{T-1} \tau_{t} e^{-r_{p}t} + (\Delta_{i} - \sigma_{T}) e^{-r_{p}T} + \sum_{t=T+1}^{\infty} \left\{ \int_{-\infty}^{\tau_{t} - \Delta_{i}} (\Delta_{i} + \delta) dF(\delta) + \left[1 - F(\tau_{t} - \Delta_{i})\right] \tau_{t} \right\} e^{-r_{p}t}$$
(15)

We explicitly allow the private discount rate to differ from the social discount rate. There is a large literature on private and social discount rates, with some arguing that they should be the same (e.g. Baumol, 1968) and others indicating that investment has a public good component and should therefore be rewarded by a social discount rate that is below the private one (Weitzman, 1994).

The first term in (15) describes the costs before learning when the consumer uses the dirty alternative, which simply consist of paying the tax τ_t in every period up to time T-1, discounted to present value by the interest rate r_p . The second term reflects the (discounted) expected cost of learning at time T, reduced by the first-time user subsidy σ_t . The last term represents costs that accrue after the consumer has learnt her value for δ_i . The integral represents the expected cost of remaining in alternative B, conditional on δ_i turning out to be low enough for this to be optimal (i.e. $\theta_i = \delta_i + \Delta_i \leq \tau_t$). Conversely, if the consumer's private costs turn out to be above this threshold she will revert to the dirty alternative and once again incur the instantaneous cost τ_t , possibly switching to the clean alternative again once the tax has increased to a sufficiently high level. This leads us to our second result:

Proposition 2

The first-best transition path can be implemented using a usage tax τ_t for the dirty alternative in combination with a subsidy for first-time users of σ_t .

Proof: In order for all agents who have already learnt to face the same incentive to use the clean alternative as in the first-best transition path (θ_t, s_t, k_t) , the regulator needs to set $\tau_t = \theta_t$. The building of new capacity k_t can directly be controlled. For generating the optimal incentives for agents to learn in the respective periods, we need that type Δ_t as defined in the optimal approach path is indifferent between learning in period t-1 and t. Considering (15), this is the case if

$$\Delta_{t} - \sigma_{t-1} + e^{-r_{p}} \left\{ \int_{-\infty}^{\tau_{t} - \Delta_{t}} (\Delta_{t} + \delta) dF(\delta) + \left[1 - F(\tau_{t} - \Delta_{t}) \right] \tau_{t} \right\} = \tau_{t-1} + (\Delta_{t} - \sigma_{t}) e^{-r_{p}}$$
(16)

It is instructive to generate more insights into the properties of this subsidy path for first-time users. In the Appendix, we show that

We can summarize this in the following result:

Proposition 3

If social and private discount rates coincide, the first-best solution can be decentralized by taxing the dirty alternative, whereas no subsidy for first-time users is necessary. If the private discount rate is larger (smaller) than the social discount rate, a subsidy (tax) for first-time users is necessary to implement the first-best solution. The discounted absolute value of this subsidy (tax) decreases over time, i.e. $|\sigma_{t-1}| > |\sigma_t| \cdot e^{-r_p}$.

Proposition 3 demonstrates that it is relatively easy to decentralize the first-best solution as long as individuals correctly weigh their future costs and benefits from trying the new alternative. The reason is that our model assumes efficient learning. If, however, the private and the social discount rates differ, relying on a tax for using the dirty alternative is not sufficient. If the private discount rate is larger, individuals do not have sufficient incentives to learn. Suboptimal learning by a consumer creates an indirect externality to other consumers: as they do not learn and do not use the cleaner alternative B, the marginal social damages from the dirty alternative and, as a result, the environmental tax is inefficiently high, thereby preventing some consumers from trying the clean alternative even if this were socially optimal. Consumers do not consider this pecuniary externality when solving the optimal stopping problem. Incentives for efficient learning can be

established by subsidizing first-time users (resp. a tax if $r_p < r$), and the absolute value of the discounted subsidy must decrease over time. Proposition 3 thereby indicates a reason for why regulators may adjust their policies in a dynamic setting, specifically temporarily subsidize trying a new alternative: while taxes are increasing over time, the subsidy for first-time users decreases.

Similarly, consumers may have biased expectations in the sense that their perceived distribution of δ differs from the true one such that $\tilde{F}(\delta) \neq F(\delta)$, i.e. if they over- or underestimate the personal costs. , we show in the Appendix that

If one distribution first-order stochastically dominates the other, b_t is always either positive or negative dominance. We immediately obtain the following corollary to Proposition 3:

Corollary

If consumers have overestimate costs δ such that \tilde{F} first-order stochastically dominates F ($\tilde{F}(\delta) < F(\delta)$ for all δ), a subsidy for first-time users is necessary to implement the first-best solution. If they underestimate the costs such that F first-order stochastically dominates \tilde{F} , a tax for first-time users is optimal. The absolute value of the discounted subsidy or tax, respectively, decreases over time, i.e. $|\sigma_{t-1}| > |\sigma_t| \cdot e^{-r}$. To illustrate this point, consider again Figure 2. If consumers systematically over-estimate their personal costs of using the clean alternative such that $b_t > 0 \forall t$,¹⁵ the change in (11) implies an increase in $\hat{\theta}(\Delta_t)$ that is qualitatively similar to an increase of the personal rate of discount, leading to a shift in the steady state from *SS* to *SS*'.

A subsidy for first-time users allows decentralizing the first-best solution, but it may not always be a feasible policy. In the next section we therefore explore the situation where the regulator cannot differentiate between first-time and other users.

5. Decentralization with a tax on the dirty alternative only

We now use these insights to determine the path for τ_t , rendered second-best by imposing $\sigma_t = 0$. The policy-path $(\tau_t)_t$ leads consumers to choose their optimal moment of learning $T(\Delta)$ according to (15) and thus generates a learning path $(\Delta_t)_t$. The social planner's problem can therefore be stated as

$$\min_{\tau_{t},k_{t}} \sum_{t=0}^{\infty} e^{-rt} \cdot \left(\int_{-\infty}^{\lambda_{t}} \int_{-\infty}^{\tau_{t}-\Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_{t}}^{\lambda_{t+1}} \Delta dG(\Delta) + D(1 - \Omega_{t}) + C(\Omega_{t}) + A(k_{t}) \right) \\
s.t. \quad \Delta_{t+1} = \Delta_{t} + f(\tau_{t}) \\
K_{t+1} = K_{t} + k_{t}; \\
K_{t+1} \ge \Omega_{t} \\
\tau_{t} \ge 0; k_{t} \ge 0; \\
\Delta_{t=0} = \Delta_{0}; K_{t=0} = K_{0} \\
with \qquad \Omega_{t} \equiv \int_{-\infty}^{\Delta_{t}} F(\tau_{t} - \Delta) dG(\Delta) + G(\Delta_{t+1}) - G(\Delta_{t})$$
(19)

¹⁵ Stutzer et al. (2011) argue the benefits of some public programs are only fully appreciated upon reflection by consumers, which does not take place without some form of stimulus. This is equivalent to private expecations $b_t > 0$.

In order to explore the properties of the second best regulation path, we start by focusing on the learning decision. It is helpful to define a personal learning threshold by the level of a *constant* tax rate $\hat{\tau}(\Delta)$ which makes individuals of type Δ indifferent between learning and never learning. It is given by

$$\Delta + \sum_{t=1}^{\infty} \left\{ \int_{-\infty}^{\hat{\tau} - \Delta} (\Delta + \delta) dF(\delta) + (1 - F(\hat{\tau} - \Delta)) \hat{\tau} \right\} e^{-r_p t} = \hat{\tau} \sum_{t=0}^{\infty} e^{-r_p t}$$

$$\Leftrightarrow (e^{r_p} - 1)(\hat{\tau} - \Delta) = \int_{-\infty}^{\hat{\tau} - \Delta} (\Delta + \delta - \hat{\tau}) dF(\delta)$$
(20)

In general, however, the tax path is not constant. Here, the threshold $\hat{\tau}(\Delta)$ helps to characterize the timing of learning that is formulated in the following lemma (proof in Appendix):

Lemma 1.

For a given path of usage taxes (τ_t) , individual of type Δ will either decide to try the clean alternative in a period t with $\hat{\tau}(\Delta) < \tau_t$ or never.

In particular, Lemma 1 implies that agents will not learn along a decreasing portion of the taxation path. Furthermore agents of type Δ will never learn if the tax rate never exceeds $\hat{\tau}(\Delta)$. They will only consider to learn in a period where $\tau_{t-1} < \hat{\tau}(\Delta) \le \tau_t$.

We use these insights to determine the second-best tax path τ_t . If no more learning takes place $(\Delta_{t+1} = \Delta_t)$ and the capacity is built up sufficiently, the optimization necessarily implies $\tau^{SS} = D'(1 - \Omega^{SS}) - C'(\Omega^{SS})$, i.e. $\tau^{SS} = \tilde{\theta}(\Delta^{SS})$ as defined in (13). That is, the relationship between the optimal tax and net marginal damages remains unchanged. Since at this tax rate no uninformed agent can have an incentive to learn we further know that $\theta^{SS} \leq \hat{\tau}(\Delta^{SS})$.

If we were to rely on a continuous and increasing approach path for the taxes, this implies that the steady state satisfies $\theta^{SS} = \hat{\tau}(\Delta^{SS}) = \tilde{\theta}(\Delta^{SS})$. We denote this private steady state by $(\theta_p^{SS}, \Delta_p^{SS})$.

In the following we will discuss the implications of differing discount rates for our model. It is clear that the first-best path can be obtained using a tax only if $r = r_s$: in this case $\hat{\tau}(\Delta) = \hat{\theta}(\Delta)$ such that the social steady state $(\theta_s^{SS}, \Delta_s^{SS})$ given by $\theta^{SS} = \hat{\theta}(\Delta^{SS}) = \tilde{\theta}(\Delta^{SS})$ and the private steady state $(\theta_p^{SS}, \Delta_p^{SS})$ coincide. However, the decentralization leads to different learning behavior if $r_p > r$, in which case we have $\hat{\tau}(\Delta) > \hat{\theta}(\Delta)$. Since $\tilde{\theta}(\Delta)$ is decreasing in Δ , this implies that

$$\theta_s^{SS} < \theta_P^{SS} \qquad \text{and} \qquad \Delta_s^{SS} > \Delta_P^{SS}.$$
(21)

Intuitively, private impatience makes people hesitant to learn. In order to induce the same consumer type to learn, a larger tax rate is needed, leading to a steady state in which fewer consumers are informed, and consequently net marginal damages and the tax/subsidy are higher.

When relying on the usage tax alone, any increasing tax path in the decentralized setting must optimally end up in the steady state $(\theta_p^{SS}, \Delta_p^{SS})$. However, we show that one can improve upon this by using a non-monotonic path.

Proposition 4.

If the private discount rate r_p exceeds the social discount rate r, a taxation path that is first increasing, but will decrease at one point in time before being constant, can improve upon a policy that relies on a monotonically increasing tax path.

We prove this proposition by noting that a steady state requires that no additional learning takes place. Without trying to trigger additional learning, the welfare-maximizing policy for any assumed steady state learning rate Δ^{∞} is given by $\tau^{\infty} = \tilde{\theta}(\Delta^{\infty})$. Assuming welfare-maximization by

the regulator, no increasing tax path can hence improve upon the permanent steady state $(\theta_p^{SS}, \Delta_p^{SS})$ once the learning rate of Δ_p^{SS} has been obtained.¹⁶

The regulator could, however, temporarily increase the tax to $\tau_t > \tilde{\theta}(\Delta_t)$ in order to induce more agents to learn (such that $\Delta_p > \Delta_p^{SS}$), before adjusting the tax to a lower level $\tau_p = \tilde{\theta}(\Delta_p) < \theta_p^{SS}$. In the Appendix, we show that this is indeed optimal with the simplest example of a path that involves a tax rate $\tau_T > \theta_p^{SS}$ for only one period before permanently being adjusted to a new lower level $\tau_p = \tilde{\theta}(\Delta_p)$.

Proposition 4 thereby complements the result stated in Proposition 3: if the private discount rate exceeds the social one, the optimal policy involves either a subsidy for first-time users that may decrease over time or a tax path that is not monotonic, but rather consists of an increasing part before being reduced to a lower level. The intuition is that one needs to induce agents to learn by trying the new clean alternative. If they have done so to a sufficient amount, the relative price of the dirty alternative can be reduced. Again, the same conclusions apply for biased expectations of costs or benefits.

6. Conclusions

Regulation often involves introductory taxes or subsidies that may later be reduced. Usually this is motivated by supply-side considerations such as decreasing production costs or technology spillovers. In this article, we discuss a different rationale for introductory subsidies that is motivated by the demand-side: If consumers are uncertain about their tastes regarding a new product, they may learn by trying out the new technology. A prominent and environmentally highly relevant example involves the introduction of new public transport options. While

¹⁶ Any announcement of further increasing tax rates to trigger learning beyond $\Delta_p SS$ would not correspond to a timeconsistent policy as the regulator would have an incentive to lower the tax rates below $\tilde{\theta}(\Delta_p^{SS})$ once the additional learning has occurred.

consumers are usually experienced in their status quo alternative (private car use), at least a part of the true opportunity costs of using public transport need to be experienced before they are known.

In this paper, we demonstrate how regulation should incorporate dynamic features that initiate from this "learning-by-trying". Any regulation needs to simultaneously account for two dimensions: experienced consumers will use the new technology if their private opportunity costs are outweighed by tax on the old technology. Second, the policy in its introductory phase needs to control the optimal number of *new* consumers.

We demonstrate that the optimal transition path into a steady state involves increasing regulation levels as long as social and private learning incentives coincide. In this case, the firstbest path can be decentralized by taxing the dirty alternative. Along the optimal path, the tax rates are increasing, which corresponds to a slow introduction of taxes. With zero or low fixed costs to expand the capacity of the new technology, the fraction of consumers using the clean alternative actually falls over time since initially many agents try out the alternative, but then may (temporarily) go back to the polluting option. If the capacity expansion is costly, both the optimal tax and usage rates increase over time.

The qualitative features of the optimal policy significantly change if the private rate of discount exceeds the social discount rate. Due to the divergence between private and social learning incentives, the first-best policy involves a subsidy for first-time users. This subsidy is decreasing over time. If such a special treatment of first-time users is not feasible, the regulator's second best tax path also needs to take into account the different learning incentives. The optimal tax path is necessarily first increasing, until the (second-best) optimal number of agents has learnt, before being adjusted downwards to the steady state, implying an initially increasing and then decreasing optimal tax.

Although we derive our results in the context of differing private and social discount rates, they also apply to situations where consumers have biased expectations about the full costs or benefits of a new technology. Like private discount rates that differ from social ones, biased expectations lead to socially inefficient learning that has to be corrected by means of subsidies/taxes for first-time users to obtain the first-best solution (where the subsidy is equal to the bias), or by a non-monotonic regulation path if the first-best solution is not feasible.

More generally, our results suggest that if private learning incentives lead to a rate of exposure to a new experience good that lies below the social optimum, introductory subsidies can be justified not only with decreasing production costs, but also when consumers learn by experience. We believe that exploring further behavioral rationales for dynamic adjustments in regulation levels provides a valuable path for future research.

References

Ackerberg, Daniel A. (2003). "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination." *International Economic Review* 44: 1007–40.

Akerlof, George A. (1970). "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism". *Quarterly Journal of Economics* 84 (3): 488–500

Andreoni, James (1990). "Impure Altruism and Donations to Public Goods: A Theory of Warm-Glow Giving". *The Economic Journal* 100: 464-477.

Baumol, William J. (1968). "On the Social Rate of Discount". *American Economic Review* 58 (4): 788-802.

Benabou, Roland and Jean Tirole (2006). "Incentives and Prosocial Behavior," *American Economic Review* 96: 1652-1678.

Bergemann, Dirk and Juuso Välimäki (2006), "Dynamic Pricing of new experience goods," *Journal of Political Economy* 114(4): 713-743.

Crawford, Gregory S. and Matthew Shum (2005). "Uncertainty and Learning in Pharmaceutical Demand." *Econometrica* 73 (July): 1137–73.

Cremer, Jacques (1984). "On the Economics of Repeat Buying." *Rand Journal of Economics* 15: 396–403.

Erdem, Tulin, Susumu Imai, and Michael P. Keane (2003). "Brand and Quantity Choice Dynamics under Price Uncertainty." *Quantitative Marketing and Economics* 1 (1): 5–64.

Erdem, Tulin, and Michael P. Keane (1996). "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets." *Marketing Science* 15 (1): 1–20.

Farrell, Joseph (1986). "Moral Hazard as an Entry Barrier." *Rand Journal of Economics* 17: 440–49.

Goettler, Ronald L., and Karen Clay (2006). "Price Discrimination with Experience Goods: A Structural Econometric Analysis." *Manuscript, Carnegie Mellon Univ*.

Israel, Mark (2005). "Service as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance." *American Economic Review* 95: 1444–63.

Milgrom, Paul, and John Roberts (1986). "Price and Advertising Signals of Product Quality." *Journal of Political Economy* 94: 796–821.

Osborne, Martin (2005). "Consumer Learning, Habit Formation and Heterogeneity: A Structural Examination." *Manuscript, Univ. Chicago*.

Shapiro, Carl (1983). "Optimal Pricing of Experience Goods." *Bell J. Econ.* 14 (Autumn): 497–507.

Stutzer, Alois, Lorenz Goette and Michael Zehnder. "Active Decisions and Prosocial Behavior: A Field Experiment on Blood Donation." *Economic Journal* 121(556): F476-F493.

Tirole, Jean (1988). The Theory of Industrial Organization. Cambridge, MA: MIT Press.

Weitzman, M. (1994). "On the "Environmental" Discount Rate." *Journal of Environmental Economics and Management* 26, 200-209.

Figure 1: Phase diagram for an interior solution in the first-best policy case



Figure 2: Change of steady state in response to underlying parameters



















Fig. 3e: Rate of expansion

Fig. 3f: Usage rate

Appendix

Derivation of eqs. (10)-(11)

For a continuous approach path, (6) will hold with equality during the approach path as well as in the steady state. Solving for $J_{\Delta}(\Delta_{t+1}, K_{t+1})$ while using (5) leads to

$$\frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial \Delta} = e^r (\theta_t - \Delta_{t+1}) g(\Delta_{t+1})$$
(A.1)

Shifting the equation by one period yields an expression for $J_{\Delta}(\Delta_t, K_t)$. Substituting both shadow prices into (9), canceling terms and dividing by $g(\Delta_t)$:

$$e^{r}(\theta_{t} - \Delta_{t+1}) = \int_{-\infty}^{\theta_{t} - \Delta_{t}} (\Delta_{t} + \delta) dF(\delta) - \theta_{t} F(\theta_{t} - \Delta_{t}) - \Delta_{t} + \theta_{t}$$
(A.2)

Including $\theta_t F(\theta_t - \Delta_t)$ in the integral and solving for $e^r(\theta_t - \theta_{t-1})$ leads to (11). Now we solve (6) for $J_K(\Delta_{t+1}, K_{t+1})$:

$$\frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial K} = e^r (\lambda_t - A_t)$$
(A.3)

Deriving $J_K(\Delta_t, K_t)$ by shifting one period and substituting both into (10) and rearranging and shifting the time period one forward leads to (12).

Derivation of eqs. (13)-(14)

Setting $s_t = 0$ and holding λ_t constant in eq. (5) gives

$$\theta_t \Big|_{s_t=0; \lambda_t=\overline{\lambda}} = D' \Big[1 - \Omega_t \Big|_{s_t=0} \Big] - C' \Big[\Omega_t \Big|_{s_t=0} \Big] - \overline{\lambda}$$

$$= D' \Big[1 - \int_{-\infty}^{\Lambda_t} F(\theta_t - \Delta) dG(\Delta) \Big] - C' \Big[\int_{-\infty}^{\Lambda_t} F(\theta_t - \Delta) dG(\Delta) \Big] - \overline{\lambda}$$
(A.4)

Totally differentiating and rearranging shows that the relationship $\tilde{\theta}(\Delta_t)$ is decreasing:

$$d\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}} = \frac{\partial\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}}}{\partial\Delta_{t}} d\Delta_{t} + \frac{\partial\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}}}{\partial\theta_{t}} d\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}}$$

$$\frac{d\tilde{\theta}(\Delta_{t})}{d\Delta_{t}} = \frac{d\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}}}{d\Delta_{t}} = \frac{\partial\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}} / \partial\Delta_{t}}{1 - \partial\theta_{t}\Big|_{s_{t}=0;\lambda_{t}=\bar{\lambda}} / \partial\theta_{t}}$$

$$= \frac{-(D^{"}+C^{"})F(\theta_{t}-\Delta_{t})g(\Delta_{t})}{1 + (D^{"}+C^{"})\int_{-\infty}^{\Delta_{t}} f(\theta_{t}-\Delta)dG(\Delta)} < 0$$
(13)

The relationship is negative due to our assumption that D">0 and C">0. Next, setting $\theta_t = \theta_{t-1}$ in (11) and solving for θ_t leads to

$$\theta_t \Big|_{\theta_t = \theta_{t-1}} = \Delta_t + \frac{1}{1 - e^r} \int_{-\infty}^{\theta_t - \Delta_t} (\theta_t - \Delta_t - \delta) dF(\delta)$$
(A.5)

Totally differentiating and rearranging gives:

$$\frac{d\hat{\theta}(\Delta_{t})}{d\Delta_{t}} = \frac{d\theta_{t}\big|_{\theta_{t}=\theta_{t-1}}}{d\Delta_{t}} = \frac{\partial\theta_{t}\big|_{\theta_{t}=\theta_{t-1}}/\partial\Delta_{t}}{1-\partial\theta_{t}\big|_{\theta_{t}=\theta_{t-1}}/\partial\theta_{t}} = \frac{1-e^{r}-\int_{-\infty}^{\theta_{t}-\Delta_{t}}dF(\delta)}{1-e^{r}-\int_{-\infty}^{\theta_{t}-\Delta_{t}}dF(\delta)} = 1$$
(14)

Derivation of eqs. (17) and (18)

Rearranging (16) leads to

$$\sigma_{t-1}e^{r_p} - \sigma_t = e^{r_p}(\Delta_t - \tau_{t-1}) + (\tau_t - \Delta_t) + \int_{-\infty}^{\tau_t - \Delta_t} (\Delta_t + \delta - \tau_t) dF(\delta)$$
(A.6)

Eq. (11) can be rearranged to

$$e^{r}(\Delta_{t} - \theta_{t-1}) + (\theta_{t} - \Delta_{t}) + \int_{-\infty}^{\theta_{t} - \Delta_{t}} (\Delta_{t} + \delta - \theta_{t}) dF(\delta) = 0$$
(A.7)

When substituting $\theta_t = \tau_t$, the RHS in (A.6) is identical to the LHS in (A.7) except for the interest rate. Whether it is greater or smaller depends on the sign of $(\Delta_t - \theta_{t-1})$. Note that (A.1)

when shifted by one period implies that $\Delta_t - \theta_{t-1} > 0$ as $J_{\Delta}(\Delta_t, K_t) \le 0$ (the presence of more informed consumers can only lower the social costs).

For $r_p > r$, combining (A.6) and (A.7) leads to

$$\sigma_{t-1}e^{r_{p}} - \sigma_{t}$$

$$= e^{r_{p}} \left(\underbrace{\Delta_{t} - \theta_{t-1}}_{>0} \right) + \left(\theta_{t} - \Delta_{t}\right) + \int_{-\infty}^{\theta_{t} - \Delta_{t}} \left(\Delta_{t} + \delta - \theta_{t}\right) dF(\delta)$$

$$(A.8)$$

$$= 0$$

$$(A.8)$$

For $r_p < r$, the inequality is reversed, and for $r_p = r$ (A.8) holds with equality, as in (17).

To derive (18), we assume that consumers use the socially optimal discount rate, but that they have biased expectations in the sense their perceived p.d.f. for δ differs from the true p.d.f. such that $\tilde{F}(\delta) \neq F(\delta)$. Defining consumers' bias as

$$b_{t} \equiv \int_{-\infty}^{\theta_{t} - \Delta_{t}} (\Delta_{t} - \theta_{t} + \delta) \tilde{F}(\delta) - \int_{-\infty}^{\theta_{t} - \Delta_{t}} (\Delta_{t} - \theta_{t} + \delta) F(\delta)$$
(A.9)

the individual learning condition that corresponds to (A.6) becomes (setting $r_p = r$, and replacing $F(\delta)$ by $\tilde{F}(\delta)$):

$$\sigma_{t-1}e^r - \sigma_t = e^r(\Delta_t - \tau_{t-1}) + (\tau_t - \Delta_t) + \int_{-\infty}^{\tau_t - \Delta_t} (\Delta_t + \delta - \tau_t) dF(\delta) + b_t$$
(A.10)

Substituting $\theta_t = \tau_t$ and using (A.7) we obtain (18).

Proof of Lemma 1.

The proof relies on showing that $\tau_t < \hat{\tau}(\Delta)$ implies that learning in period *t* cannot be optimal for type Δ .

(i) We first show that showing this is also sufficient to rule out the case where $\tau_t < \tau_{t-1}$. If learning in period *t* leads to smaller costs than learning in period *t*-1, we obtain

$$\Delta + e^{-r_p} \left\{ \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta) dF(\delta) + (1 - F(\tau_t - \Delta))\tau_t \right\} \ge \tau_{t-1} + e^{-r_p} \Delta$$

$$\Leftrightarrow e^{r_p} (\tau_t - \tau_{t-1}) \ge (e^{r_p} - 1)(\tau_t - \Delta) + \int_{-\infty}^{\tau_t - \Delta} (\tau_t - \Delta - \delta) dF(\delta)$$
(A.11)

which would imply $\tau_t < \hat{\tau}(\Delta)$.

(ii) Now assume that $\tau_t < \hat{\tau}(\Delta)$ and that tax rates $\tau_s < \hat{\tau}(\Delta)$ for all future periods s > t. Here, never learning generates less costs than learning in t for type Δ since with the definition of $\hat{\tau}(\Delta)$ in (20), we obtain:

$$\begin{split} &\left[\Delta + \sum_{s=t+1}^{\infty} \left\{\int^{\tau_s - \Delta} (\Delta + \delta) dF(\delta) + (1 - F(\tau_s - \Delta))\tau_s\right\} e^{-r_p(s-t)} \right] - \left[\sum_{s=t}^{\infty} \tau_s e^{-r_p(s-t)} \right] \\ &= \Delta - \tau_t + \sum_{s=t+1}^{\infty} \left\{\int^{\tau_s - \Delta} (\Delta + \delta - \tau_s) dF(\delta)\right\} e^{-r_p(s-t)} \\ &\geq \Delta - \hat{\tau}(\Delta) + \sum_{s=t+1}^{\infty} \left\{\int^{\hat{\tau}(\Delta) - \Delta} (\Delta + \delta - \hat{\tau}(\Delta)) dF(\delta)\right\} e^{-r_p(s-t)} \\ &= 0 \end{split}$$

(iii) It remains the case where there exists a period where $\tau_T > \tau_t$. Define *T* as the first period in which $\tau_T > \tau_t$, i.e. $\tau_s \le \tau_t$ for all $t \le s < T$. We show that costs from learning in *t* are larger than learning in *T*:

$$\begin{split} \left[\Delta + \sum_{s=t+1}^{T} \left\{ \int_{-\infty}^{\tau_{s}-\Delta} (\Delta + \delta) dF(\delta) + (1 - F(\tau_{s} - \Delta))\tau_{s} \right\} e^{-r_{p}(s-t)} \right] - \left[\tau_{t} + \sum_{s=t+1}^{T-1} \tau_{s} e^{-r_{p}(s-t)} + \Delta e^{-r_{p}(T-t)} \right] \\ &= \Delta - \tau_{t} + \sum_{s=t+1}^{T-1} \left\{ \int_{-\infty}^{\tau_{s}-\Delta} (\Delta + \delta - \tau_{s}) dF(\delta) \right\} e^{-r_{p}(s-t)} + \left(\int_{-\infty}^{\tau_{p}-\Delta} (\Delta + \delta - \tau_{T}) dF(\delta) - (\Delta - \tau_{T}) \right) e^{-r_{p}(T-t)} \\ &> \Delta - \tau_{t} + \int_{-\infty}^{\tau_{t}-\Delta} (\Delta + \delta - \tau_{t}) dF(\delta) \sum_{s=t+1}^{T-1} e^{-r_{p}(s-t)} + \left(\int_{-\infty}^{\tau_{p}-\Delta} (\Delta + \delta - \tau_{t}) dF(\delta) - (\Delta - \tau_{t}) \right) e^{-r_{p}(T-t)} \\ &= (\Delta - \tau_{t}) (1 - e^{-r_{p}(T-t)}) + \int_{-\infty}^{\tau_{t}-\Delta} (\Delta + \delta - \tau_{t}) dF(\delta) \sum_{s=t+1}^{T} e^{-r_{p}(s-t)} \\ &> \left[\Delta - \tau_{t} \right] \left[1 - (e^{r_{p}} - 1) \sum_{s=t+1}^{T} e^{-r_{p}(s-t)} - e^{-r_{p}(T-t)} \right] \\ &= \left[\Delta - \tau_{t} \right] \left[1 - (e^{r_{p}} - 1) e^{-r_{p}} \frac{1 - e^{-r_{p}(T-t)}}{1 - e^{-r_{p}}} - e^{-r_{p}(T-t)} \right] \\ &= 0 \end{split}$$

Here, we used $\tau_T > \tau_t$ and $\tau_s \le \tau_t$ for all $t \le s < T$, to obtain the first inequality and $\tau_t < \hat{\tau}(\Delta)$ for the second inequality.

Summarizing (ii) and (iii) shows that an agent of type Δ may only decide to learn if $\tau_t > \hat{\tau}(\Delta)$. With (i), he may learn in period *t* if $\tau_t > \tau_{t-1}$, i.e. along an increasing portion of the tax path.

Proof of Proposition 4.

We assume that the current $\Delta_T = \Delta_P^{SS}$. In order to induce a marginal Δ_P -type consumer to learn by just increasing the tax rate for one period before reducing it to $\tilde{\theta}(\Delta_P)$ in such a setting, the tax rate τ_T needs to satisfy the following learning condition:

$$\Delta_{p} + \frac{e^{-r_{p}}}{1 - e^{-r_{p}}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_{p}) - \Delta_{p}} (\Delta_{p} + \delta) dF(\delta) + (1 - F(\tilde{\theta}(\Delta_{p}) - \Delta_{p}))\tilde{\theta}(\Delta_{p}) \right\} = \tau_{T} + \frac{e^{-r_{p}}}{1 - e^{-r_{p}}} \tilde{\theta}(\Delta_{p}) \quad (A.13)$$

which can be solved for

$$\tau_T = \tau_T(\Delta_P) = \Delta_P + \frac{e^{-r_P}}{1 - e^{-r_P}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_P) - \Delta_P} (\Delta_P + \delta - \tilde{\theta}(\Delta_P)) dF(\delta) \right\}$$
(A.14)

The social costs from this path are given by:

$$\operatorname{Cost}(\Delta_{p}) = \left(\int_{-\infty}^{\Delta_{p}^{SS}} \int_{-\infty}^{\tau_{T}(\Delta_{p})-\Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_{p}^{SS}}^{\Delta_{p}} \Delta dG(\Delta) + D[1-\Omega_{T}] + C[\Omega_{T}] + A(\Omega_{T} - K_{T})\right) + \frac{e^{-r}}{1 - e^{-r}} \left(\int_{-\infty}^{\Delta_{p}} \int_{-\infty}^{\tilde{\theta}(\Delta_{p})-\Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + D[1-\Omega_{p}] + C[\Omega_{p}]\right)$$
(A.15)

with

$$\Omega_{T} = \int^{\Delta_{P}^{SS}} F(\tau_{T}(\Delta_{P}) - \Delta) dG(\Delta) + (G(\Delta_{P}) - G(\Delta_{P}^{SS}))$$

$$(A.16)$$

$$\Omega_{P} = \int^{\Delta_{P}} F(\tilde{\theta}(\Delta_{P}) - \Delta) dG(\Delta)$$

It is sufficient to show that $\operatorname{Cost}'(\Delta_p^{SS}) < 0$. Noting that $\theta_p^{SS} = \tau_T(\Delta_p) = \tilde{\theta}(\Delta_p)$ in $\Delta_p = \Delta_p^{SS}$, we obtain:

$$\begin{aligned} \operatorname{Cost}'(\Delta_{P}^{SS}) &/ g(\Delta_{P}^{SS}) \\ &= \Delta_{P}^{SS} - (\underbrace{D'(1 - \Omega_{T}) - C'(\Omega_{T})}_{=\tilde{\theta}(\Delta_{P}^{SS})}) \\ &+ \frac{e^{-r}}{1 - e^{-r}} \begin{cases} \tilde{\theta}(\Delta_{P}^{SS}) - \Delta_{P}^{SS}}{\int_{-\infty}^{\infty}} (\Delta_{P}^{SS} + \delta) dF(\delta) - F(\tilde{\theta}(\Delta_{P}^{SS}) - \Delta_{P}^{SS})(\underbrace{D'(1 - \Omega_{P}) - C'(\Omega_{P})}_{=\tilde{\theta}(\Delta_{P}^{SS})})) \end{cases} \\ &= \Delta_{P}^{SS} - \tilde{\theta}(\Delta_{P}^{SS}) + \frac{e^{-r}}{1 - e^{-r}} \begin{cases} \tilde{\theta}(\Delta_{P}^{SS}) - \Delta_{P}^{SS}}{\int_{-\infty}^{\infty}} (\Delta_{P}^{SS} + \delta - \tilde{\theta}(\Delta_{P}^{SS})) dF(\delta) \end{cases} \\ &< \Delta_{P}^{SS} - \tilde{\theta}(\Delta_{P}^{SS}) + \frac{e^{-r_{P}}}{1 - e^{-r_{P}}} \begin{cases} \tilde{\theta}(\Delta_{P}^{SS}) - \Delta_{P}^{SS}}{\int_{-\infty}^{\infty}} (\Delta_{P}^{SS} + \delta - \tilde{\theta}(\Delta_{P}^{SS})) dF(\delta) \end{cases} \\ &= 0 \end{aligned} \end{aligned}$$

$$(A.17)$$

where the last inequality follows from the assumed relationship between private and social discount rates $r_p > r$. The last equality follows from the learning condition in the private steady state.