

Learning abatement costs: On the dynamics of the 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 on the polluting consumption is increasing over time. We show, however, that if social and private learning incentives diverge because the private discount rate exceeds the social discount rate, it may be optimal to temporarily increase the tax rate beyond net marginal external damages to induce more learning before reducing the tax rate to the steady-state level. Alternatively, one could complement the tax with subsidies for first-time users which can be phased out over time. Similar results apply if consumers have biased expectations. We therefore give a rationale for introductory subsidies on new, clean technologies and non-monotonic tax paths from a perspective of consumer learning.

JEL: D82, L51, H21, H23, Q52, Q55

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

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 the personal net costs they incur from its use, i.e. their personal abatement costs.

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 advertised in terms of their higher energy efficiency relative to the older versions they seek to replace, but consumers only have limited knowledge about the operating costs and convenience associated with these new products. Learning by experience can also be linked to motivational factors such as the warm-glow effect or social reputation derived from using environmentally friendly products (Andreoni 1990, Benabou and Tirole 2006). These examples share two features: 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, explore rationales for initial subsidies for a new product, and demonstrate how optimal regulation levels change over time. We focus on the dynamics of a government intervention that is driven exclusively by consumers learning the net costs that they personally incur from adopting a new technology (hereafter denominated ‘learning’). That is, we abstract from learning spillovers among consumers as well as from supply-side

arguments such as cost reductions through an intensified use of the new technology 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 the supply of experience goods under imperfect competition. Bergemann and Välimäki (2006) examine monopolistic pricing of experience goods in a dynamic model. They show that price dynamics crucially depend on a simple dichotomy between mass and niche markets.¹ While prices in mass markets are declining over time, they may initially be low but subsequently increase in niche markets. The low initial prices are set with a policy focus aimed at 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. The pricing of experience goods has also been discussed by Cremer (1984), Farrell (1986), and Milgrom and Roberts (1986).

We set up a dynamic model of a new experience good that replaces a preexisting consumption choice which generates 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

¹ In their paper, a mass market is characterized by the optimal price for informed consumers being low enough that uninformed consumers have an incentive to try the new product. In such a market, all consumers will eventually try the new product. Conversely, a niche market will feature consumers who do not try the new product, as they have no incentive to do so at the optimal price for informed consumers.

component. The latter is subject to idiosyncratic uncertainty and can only be learned through individual experience. We assume that consumers only learn their true costs, i.e. their valuation relative to the old technology, once they have used 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 who are exposed to the new technology for the first time, and the set of experienced consumers whose continued use of the technology is optimal. Second, we consider a setting where the regulator needs to rely solely on subsidies or taxes. The subsidy in the given period then 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 the net costs of the new technology are unbiased, the first-best case can be decentralized by simply relying on a corrective tax levied on the polluting technology, with the optimal tax rate increasing over time. If the private discount rate exceeds the social discount rate or if consumers' cost expectations are biased, the tax alone does not suffice. The first-best solution can then be implemented by complementing the tax with a subsidy for first-time users. This subsidy will be non-increasing over time. If the regulator cannot discriminate between first-time and experienced users, the second-best taxation scheme may involve a non-

² This is a stylized assumption. In general, learning takes places at different rates, and consumers may have different levels of 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), Crawford and Shum (2005), and Goettler and Clay (2006)).

monotonic path: Tax rates are initially increasing with a policy focused on reaching the optimal amount of knowledge (i.e. the number of consumers that have learned their net personal costs), 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 the 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 so as 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 that it may also 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 vary by individual: In other words, they do not know the true opportunity costs of adopting the new technology. Once a consumer tries B for one period, he learns his true net costs (or equivalently, his overall utility). If these net costs are below those associated with alternative A, the consumer will continue using alternative B; otherwise he will revert to alternative A.⁴ Consumers ignore environmental externalities when choosing between A and B.

We separate the private net costs to the consumer into an ‘unknown’ part prior to trying alternative B, and an ‘ex ante known’ part. 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 $\Delta \in S_\Delta$, which are distributed within the population according to the pdf $g(\Delta)$ and cdf $G(\Delta)$ defined over an interval S_Δ . In addition, each consumer faces a cost component $\delta \in S_\delta$ that is unknown before experiencing the good. We denote the pdf and cdf of δ over the interval S_δ by $f(\delta)$ and $F(\delta)$, respectively. We assume

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 who switches from A to B. 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 their car when trying the public transport option rather than selling it immediately. Other examples could include renting zero-energy housing or using 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.

that Δ and δ are distributed independently of each other and have full support. Consumers are thus defined over a type-space $\Theta \equiv S_{\Delta} + S_{\delta}$.

Private net costs are given by $\theta \equiv \Delta + \delta$. We normalize $E[\delta] = 0$ such that a consumer's type can be characterized by his /her (male form used hereafter for simplicity) ex-ante expected costs $E[\theta | \Delta] = E[\Delta + \delta | \Delta] = \Delta$.

Every period, a fraction of consumers try alternative B for the first time. The distribution of Δ defines the sequence according to which consumers try alternative B, with the first consumers to learn having 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 at time t .

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 his private net costs by trying alternative B chooses to use B if, and only if, $\theta \leq \theta_t$. Let $\Omega_t \in [0,1]$ describe the fraction of consumers who use B at time t . Normalizing the mass of consumers to 1, we obtain

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

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

The use of the dirty alternative A generates external damages $D = D(1 - \Omega_t)$ (with $D' > 0$, $D'' \geq 0$). The clean alternative B causes operating costs that depend on the number of users, $C = C(\Omega_t)$ (with $C' > 0$, $C'' \geq 0$). We assume that in any period, the usage rate of alternative B is capacity-constrained: $\Omega_t \leq K_t + k_t$ where K_t denotes the existing capacity from the previous period and k_t , its expansion in period t . Expanding the capacity (or quantity) of alternative B in period t to $K_{t+1} = K_t + k_t \geq K_t$ generates costs $Q(k_t)$ (with $Q' \geq 0$, $Q'(0) = 0$, $Q'' > 0$). 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 usually entail 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 this into the objective function, the intertemporal optimization problem can be written as minimizing social costs (SC) by choosing the control variables θ_t, s_t and

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

k_t at every point in time, subject to state equations for the state of knowledge and the available capacity, a capacity constraint, nonnegativity constraints, and initial conditions for the state of learning and capacity:

$$\begin{aligned}
\min_{\theta_t, s_t, k_t} \quad SC &= \sum_{t=0}^{\infty} e^{-rt} \cdot \left(\int_{-\infty}^{\Delta_t} \int_{-\infty}^{\theta_t - \Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_t}^{\Delta_{t+1}} \Delta dG(\Delta) + D(1 - \Omega_t) + C(\Omega_t) + Q(k_t) \right) \\
s. t. \quad \Delta_{t+1} &= \Delta_t + s_t \\
K_{t+1} &= K_t + k_t \\
K_{t+1} &\geq \Omega_t \\
\theta_t &\geq 0; s_t \geq 0; k_t \geq 0 \\
\Delta_{t=0} &= \Delta_0; K_{t=0} = K_0
\end{aligned} \tag{2}$$

The first term in SC measures the net costs of informed consumers who choose alternative B in 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 expansions, respectively. For each period, the social planner determines the cost cutoff value θ_t for informed consumers, the rate of first-time users s_t , and the capacity extension k_t .⁶ The Bellman equation corresponding to (2) is

⁶The per-period capacity extension k_t and the state of capacity 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,

$k_t = \max(\Omega_t - \Omega_{t-1}, 0)$. We chose to leave k_t in the model, because it facilitates exposition.

$$J(\Delta_t, K_t) = \min_{s_t, \theta_t, k_t} \left[\int_{-\infty}^{\Delta_t} \int_{-\infty}^{\theta_t - \Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_t}^{\Delta_t + s_t} \Delta dG(\Delta) + D(1 - \Omega_t) + C(\Omega_t) + Q(k_t) \right] + e^{-r} J(\Delta_t + s_t, K_t + k_t) \quad (3)$$

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

$$L = J(\Delta_t, K_t) - \lambda_t (K_t + k_t - \Omega_t) \quad (4)$$

Dropping the function arguments for damages D , operating costs C , and expansion costs Q (but adding a subscript to clarify that the arguments of these functions change over time), the optimality conditions are:

$$(D'_t - C'_t - \lambda_t) - \theta_t \geq 0; \quad \theta_t \geq 0 \quad (5)$$

$$[\Delta_{t+1} - (D'_t - C'_t - \lambda_t)] g(\Delta_{t+1}) + e^{-r} \frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial \Delta} \geq 0; \quad s_t \geq 0 \quad (6)$$

$$Q'_t - \lambda_t + e^{-r} \frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial K} \geq 0; \quad k_t \geq 0 \quad (7)$$

$$K_t + k_t - \Omega_t \geq 0; \quad \lambda_t \geq 0 \quad (8)$$

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 costs θ_t associated with alternative B must be equal to marginal net benefits $(D'_t - C'_t - \lambda_t)$ of using this alternative, where λ_t adjusts for the current-period value of a marginal capacity expansion. Equation (6) equates the net social costs of

subjecting the marginal uninformed consumer (the consumer with $\Delta = \Delta_{t+1}$) to learning in period t , (consisting of the marginal expected cost of using this alternative Δ_{t+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 $K_t > \Omega_t$, marginal capacity expansion has no value, in which case (8) implies that $\lambda_t = 0$. For $\lambda_t > 0$, equation (7) states that marginal expansion costs Q'_t have to be equal to the current-period marginal value of expansion λ_t plus the discounted marginal social value of having a larger capacity in the next period.

The intertemporal arbitrage conditions are given by

$$J_{\Delta}(\Delta_t, K_t) = \int_{-\infty}^{\theta_t - \Delta_t} (\Delta_t + \delta) dF(\delta) g(\Delta_t) + [\Delta_{t+1} g(\Delta_{t+1}) - \Delta_t g(\Delta_t)] - (D'_t - C'_t - \lambda_t) [F(\theta_t - \Delta_t) g(\Delta_t) + g(\Delta_{t+1}) - g(\Delta_t)] + e^{-r} J_{\Delta}(\Delta_{t+1}, K_{t+1}) \quad (9)$$

$$J_K(\Delta_t, K_t) = -\lambda_t + e^{-r} J_K(\Delta_{t+1}, K_{t+1}) \quad (10)$$

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 features a monotonic approach path for all variables, conditions (5)-(8) hold with equality during the approach path as well as in the steady state. This property can be used to combine Eqs. (6)-(10) (derivation in Appendix) to get

$$e^r (\theta_t - \theta_{t-1}) = (e^r - 1)(\theta_t - \Delta_t) + \int_{-\infty}^{\theta_t - \Delta_t} (\theta_t - \Delta_t - \delta) dF(\delta) \quad (11)$$

$$\lambda_t = Q'_t - e^{-r} Q'_{t+1} \quad (12)$$

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 (5)⁸ while holding λ_t constant,⁹ and the equation of motion for θ_t by setting $\theta_t = \theta_{t-1}$ in (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):

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

The economic interpretation of (13) is as follows: When no more learning takes place, the presence of more informed agents leads to a higher usage rate of the clean technology, *ceteris paribus*, and thus to lower marginal social costs. But this means that the cutoff value θ_t that fully

⁷ For $Q'[0] = 0$, Eq. (12) 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 $\theta_t = \theta_{t+1} = 0$.

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 $\lambda_t = 0$ in the steady state such that the third dimension would needlessly complicate the exposition.

internalizes external costs (for any given value of λ_t) will be lower. Conversely, if fewer people have learned, 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 knowledge, which we denote by $\hat{\theta}(\Delta_t)$. This equality 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):

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

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 remains constant at θ_t . It is constant in Δ_t due to our specification that a consumer's ex-post costs 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 in the figure. As a consequence, the optimal path to the steady state is characterized by cutoff 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 of a path to the steady state. For the first-best solution, we can summarize this result as follows:

Proposition 1

In the first-best transition path to the steady state, the marginal expected 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 $D'_t - C'_t$ decrease over time (recall that $D'_t - C'_t$ 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 $D'_t - C'_t$ will increase too, implying a declining

¹⁰ This latter point applies more generally, 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. Similarly, the initial capacity K_0 has to be below steady-state capacity in order for the capacity constraint to be binding.

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 costs turn out to be low enough ($\Delta + \delta \leq \theta_t$). Increasing net environmental costs $D'_t - C'_t$ therefore imply that the increase in informed users due to an increase in θ_t is over-compensated by the decline in first-time users.¹¹ Note also that a declining usage rate implies that too much capacity is being built, 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, *ceteris paribus*. With an increase in r , the steady state moves up the $\tilde{\theta}(\Delta_t)$ -line leading to a higher 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' .

¹¹ This follows from the definition of Ω_t : If $D'_t - C'_t$ 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)).

The function $\tilde{\theta}(\Delta_t)$ increases with marginal social damages $D'_t - C'_t$. 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 $\hat{\theta}_1(\Delta_t)$ and the steady state moves up along the $\hat{\theta}_0(\Delta_t)$ line to SS'' , which is associated with both a higher state of knowledge 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 $Q(k_t) = (\gamma / 2) \cdot k_t^2$.

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 cost threshold θ_t , the state of knowledge $\Delta_t - 0.5$ and total capacity K_t increase, whereas the rate of capacity expansion k_t decreases over time. The participation rate Ω_t decreases over time if capacity expansion is cheap (γ at or close to zero), because the decrease in an initially very high rate of learning overcompensates the increase in informed users. The higher the expansion costs,

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 $G(\Delta_t) = \Delta_t - 0.5$.

the slower are the increases of in θ_t , Δ_t and Ω_t . The rate of learning 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 privately-known net costs in period t .

Consumers choose if and when they wish to learn (i.e., to try the clean alternative for the first time), based on their private known costs Δ , their expected unknown costs δ , and the prevailing tax and subsidy rates. Conditional on the policy-path (τ_t, σ_t) and their private rate of discount r_p ,

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, the resulting state of knowledge 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. Both types of subsidy correct for a learning inefficiency: If consumers are reluctant to learn (meaning that $r_p > r$), they require a learning subsidy to learn according to (11). Conversely, if they are too eager to learn, they have to be deterred from learning a learning tax. We refer to a subsidy because it is difficult to motivate the situation where $r_p < r$, but the model does not exclude the possibility that $\sigma_t < 0$.

consumers of type Δ choose the optimal moment of learning $T(\Delta)$ that minimizes the current value of expected total costs by solving the stopping problem

$$\min_T \sum_{t=0}^{T-1} \tau_t e^{-r_p t} + (\Delta - \sigma_T) e^{-r_p T} + \sum_{t=T+1}^{\infty} \left\{ \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta) dF(\delta) + [1 - F(\tau_t - \Delta)] \tau_t \right\} e^{-r_p t} \quad (15)$$

That is, consumers of type Δ learn at a given time $T(\Delta)$ only if the costs of doing so are strictly lower than the costs of learning at any other $t \neq T(\Delta)$ and of never learning at all, in which case $T(\Delta) = \infty$.

We explicitly allow the private discount rate to differ from the social discount rate. There is a large literature on private versus social values, with some authors arguing that they should be equal (e.g. Baumol, 1968) and others maintaining that, since investment has a public-good component, it 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 consists of paying the tax τ_t in every period up to time $T-1$, discounted to the present using the interest rate r_p . The second term reflects the (discounted) expected costs 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 learned the value of his δ . The integral represents the expected costs of remaining in alternative B, conditional on δ turning out to be low enough for this to be optimal (i.e., $\delta + \Delta \leq \tau_t$). Conversely, if the consumer's private costs turn out to be above this threshold, he will revert to the dirty alternative and once again incur the instantaneous costs 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 with 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 learned their personal costs to face the same incentive to use the clean alternative according to 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. In order to generate the optimal incentives for agents to learn in the respective periods, consumer type Δ_t , as defined in the optimal approach path must be indifferent between learning in period $t-1$ and t . Considering (15), this is the case if

$$\Delta_t - \sigma_{t-1} + e^{-r_p} \left\{ \int_{-\infty}^{\tau_t - \Delta_t} (\Delta_t + \delta) dF(\delta) + [1 - F(\tau_t - \Delta_t)] \tau_t \right\} = \tau_{t-1} + (\Delta_t - \sigma_t) e^{-r_p} \quad \blacksquare \quad (16)$$

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

$$\begin{array}{ccc} > & & > \\ \sigma_{t-1} = e^{-r_p} \sigma_t & \text{and} & \sigma_t = 0 \quad \text{if} \quad r_p = r \\ < & & < \end{array} \quad (17)$$

We can summarize this in the following proposition:

Proposition 3

If social and private discount rates coincide, the first-best solution can be decentralized by taxing the dirty alternative. If the private discount rate is higher (lower) than the social discount rate, the first-best solution can be implemented by complementing the tax with a subsidy (tax) for first-time users. The discounted absolute value of this subsidy (tax) is decreasing 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 the future costs and benefits attributed to trying the new alternative. The reason is that our model assumes efficient learning. If, however, the private and the social discount rates differ, relying simply on a tax to curb usage of the dirty alternative is not sufficient. If the private discount rate is higher, individuals do not have sufficient incentives to learn. Suboptimal learning on the part of the consumer creates an indirect externality for other consumers: As they do not learn and do not use the cleaner alternative B, the marginal social damages from the dirty alternative and the associated environmental tax are inefficiently high, thereby preventing some consumers from trying the clean alternative even if that were the socially optimal choice. Consumers do not consider this pecuniary externality when solving the optimal stopping problem. Incentives for efficient learning can be established specifically by subsidizing the consumer's first use of the clean alternative where the absolute discounted value of the subsidy decreases over time. The steady state is reached with a strictly positive subsidy which could be completely phased out once the steady state is reached, i.e. when no additional consumer learns. If

$r_p < r$, first-time use needs to be taxed, with the absolute discounted value of the tax strictly decreasing over time.¹⁵

Proposition 3 indicates a reason for why regulators may adjust their policies in a dynamic setting, by temporarily subsidizing trials of 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 costs, $\tilde{F}(\delta)$, differs from the true distribution such that $\tilde{F}(\delta) \neq F(\delta)$; i.e., if they over- or underestimate their personal costs. We show in the Appendix that

$$\begin{aligned} & \begin{array}{ccc} & > & \\ \sigma_{t-1} = e^{-r} \sigma_t & \text{if } b_t = 0 & \\ & < & \end{array} & (18) \\ & \text{with } b_t \equiv \int_{-\infty}^{\theta_t - \Delta_t} (\Delta_t - \theta_t + \delta) d\tilde{F}(\delta) - \int_{-\infty}^{\theta_t - \Delta_t} (\Delta_t - \theta_t + \delta) dF(\delta) \end{aligned}$$

If one distribution first-order stochastically dominates the other (in the sense that $\tilde{F}(\delta) < (>) F(\delta)$ for all δ), then b_t is always either positive or negative. We immediately obtain the following corollary to Proposition 3:

¹⁵ The strict inequality in Eq. (17) implies a nonzero subsidy or tax in the steady state. For $r_p > r$, the subsidy could be reduced to zero once the steady state has been reached without affecting the learning decision (no consumer would have an incentive to learn either way). For $r_p < r$, the steady-state tax has to be kept indefinitely at

$$\sigma^{SS} = (e^{r_p} - e^r)(\Delta^{SS} - \theta^{SS}) / (e^{r_p} - 1).$$

Corollary

If consumers overestimate (underestimate) costs δ such that \tilde{F} first-order stochastically dominates (is first-order stochastically dominated by) F , the first-best solution can be implemented by augmenting the tax with a subsidy (tax) for first-time users. The absolute value of the discounted subsidy or tax is decreasing over time, i.e. $|\sigma_{t-1}| > |\sigma_t| \cdot e^{-r}$.

To illustrate this point, consider again Figure 2. If consumers systematically overestimate their personal costs of using the clean alternative such that $b_t > 0$ for all t .¹⁶ 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 the first-best solution to be decentralized, 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 determine the path for τ_t , rendered second-best by imposing $\sigma_t = 0$. The tax-path $TP = (\tau_t)$ leads consumers to choose their optimal moment of learning $T(\Delta)$ according to (15) and

¹⁶ Stutzer et al. (2011) argue that 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 expectations $b_t > 0$.

thus determines an equation of motion, $\Delta_{t+1} = \Delta_t + s_t(TP)$, for the amount of consumers Δ_t who have learned until a given time t . The social planner's problem becomes

$$\begin{aligned} \min_{TP=(\tau_t), (k_t)} \quad & \sum_{t=0}^{\infty} e^{-rt} \cdot \left(\int_{-\infty}^{\Delta_t} \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_t}^{\Delta_{t+1}} \Delta dG(\Delta) + D(1 - \Omega_t) + C(\Omega_t) + Q(k_t) \right) \\ \text{s.t.} \quad & \Delta_{t+1} = \Delta_t + s_t(TP) \\ & K_{t+1} = K_t + k_t \\ & K_{t+1} \geq \Omega_t \\ & \tau_t \geq 0; k_t \geq 0 \\ & \Delta_{t=0} = \Delta_0; K_{t=0} = K_0 \\ & \text{with} \quad \Omega_t \equiv \int_{-\infty}^{\Delta_t} F(\tau_t - \Delta) dG(\Delta) + G(\Delta_{t+1}) - G(\Delta_t) \end{aligned} \tag{19}$$

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.¹⁷ It is given by

¹⁷ The individual is also indifferent between learning at $t=0$, and at a later time $s>0$. To see this, suppose that (20) holds. Moving to a period $s>t$ leads to three changes in the first line: i.) the first term on the LHS has to be multiplied by $e^{-r_p t}$, ii.) the summation index on the LHS starts at $t=s+1$, and iii.) the summation index on the RHS starts at $t=s$. However, it is easy to verify that the second line will not change. If the consumer is indifferent between learning at $t=0$ and never, and between $t=s$ and never, he is also indifferent between learning at $t=0$ and $t=s$.

$$\begin{aligned} \Delta + \sum_{t=1}^{\infty} \left\{ \int_{-\infty}^{\hat{\tau}-\Delta} (\Delta + \delta) dF(\delta) + (1 - F(\hat{\tau} - \Delta)) \hat{\tau} \right\} e^{-r_p t} &= \hat{\tau} \sum_{t=0}^{\infty} e^{-r_p t} \\ \Leftrightarrow (e^{r_p} - 1)(\hat{\tau} - \Delta) &= \int_{-\infty}^{\hat{\tau}-\Delta} (\Delta + \delta - \hat{\tau}) dF(\delta) \end{aligned} \quad (20)$$

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

For any given path of usage taxes $TP = (\tau_t)$, individuals of type Δ may learn in a period t by trying the clean alternative in t only if (i) $\tau_t > \hat{\tau}(\Delta)$ and (ii) $\tau_t > \tau_{t-1}$.

In particular, the lemma 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 $\hat{\tau}(\Delta)$. We use these insights to determine the second-best tax path $TP = (\tau_t)$. 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 relative to the first-best outcome. Since at a steady state tax rate no uninformed agent can have an incentive to learn, we further know that $\theta^{SS} = \tau^{SS} \leq \hat{\tau}(\Delta^{SS})$. We define a private steady state $(\theta_p^{SS}, \Delta_p^{SS})$ as the solution of $\theta_p^{SS} = \hat{\tau}(\Delta_p^{SS}) = \tilde{\theta}(\Delta_p^{SS})$.

In the following, we will discuss the implications of differing discount rates for our model. It is clear from Proposition 3 that the first-best path can be obtained using a tax only if $r_p = r$: 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 a different learning behavior if $r_p > r$, in which case we have $\hat{\tau}(\Delta) > \hat{\theta}(\Delta)$, which follows immediately from comparing the right-hand side of (11) which defines $\hat{\theta}(\Delta)$ with the definition of $\hat{\tau}(\Delta)$ in (20). Since $\tilde{\theta}(\Delta)$ is decreasing in Δ , this implies that

$$\theta_s^{SS} < \theta_p^{SS} \quad \text{and} \quad \Delta_s^{SS} > \Delta_p^{SS} \quad (21)$$

Intuitively, private impatience makes people hesitant to learn. In order to induce the same consumer type to learn, a higher tax rate is needed, leading to a steady state in which fewer consumers are informed, and consequently net marginal damages and the usage tax 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.¹⁸

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 can indeed be welfare-improving 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 first-best solution can be implemented by complementing the tax with a subsidy for first-time users. Alternatively, if such a subsidy is infeasible, a taxation path which is at first increasing, but which decreases at a future point in time before becoming constant, can improve upon a policy that relies on a monotonically increasing tax path. The intuition is that one needs to induce agents to learn their personal net costs by trying the new clean alternative. If a sufficient number of agents have done so, the relative price of the dirty alternative can be reduced. Again, the same conclusions apply for biased expectations of costs or benefits.

¹⁸ Any announcement of further increasing tax rates to trigger learning beyond Δ_p^{SS} would not correspond to a time-consistent 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.

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 technology, they may learn by trying. A prominent and environmentally relevant example involves the introduction of new public transport options. Whereas consumers usually know the total private costs and benefits of their status quo alternative (commuting by car), at least a part of the true opportunity costs of using public transport need to be experienced before they can be known.

In this paper, we demonstrate how an optimal regulation should incorporate dynamic features that initiate from this “learning-by-trying” in an intertemporal setting. Any regulation needs to simultaneously account for two dimensions: First, experienced consumers will use the new technology if their private opportunity costs are outweighed by a 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 to a steady state involves increasing regulation levels as long as social and private learning incentives coincide. In this case, the first-best 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 consumers try 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 discount rate exceeds the social discount rate. Due to the divergence between private and social learning incentives, a tax on the dirty alternative alone cannot implement the first-best solution, but complementing the tax with a subsidy for first-time users can. This subsidy is (weakly) 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 the different learning incentives into account. A tax 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.

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. Similar to the inefficiencies resulting from the difference between private and social discount rates, biased expectations result in socially inefficient learning that has to be corrected by means of subsidies or taxes for first-time users in order 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.

References

- Akerberg, Daniel A. (2003). "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination," *International Economic Review* 44: 1007–1040.
- 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 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.
- 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.
- 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 University*.

Israel, Mark (2005). "Service as Experience Goods: An Empirical Examination of Consumer Learning in Automobile Insurance," *American Economic Review* 95: 1444–1463.

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

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

Stutzer, Alois, Lorenz Goette and Michael Zehnder (2011). "Active Decisions and Prosocial Behavior: A Field Experiment on Blood Donation," *Economic Journal* 121(556): F476-F493.

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

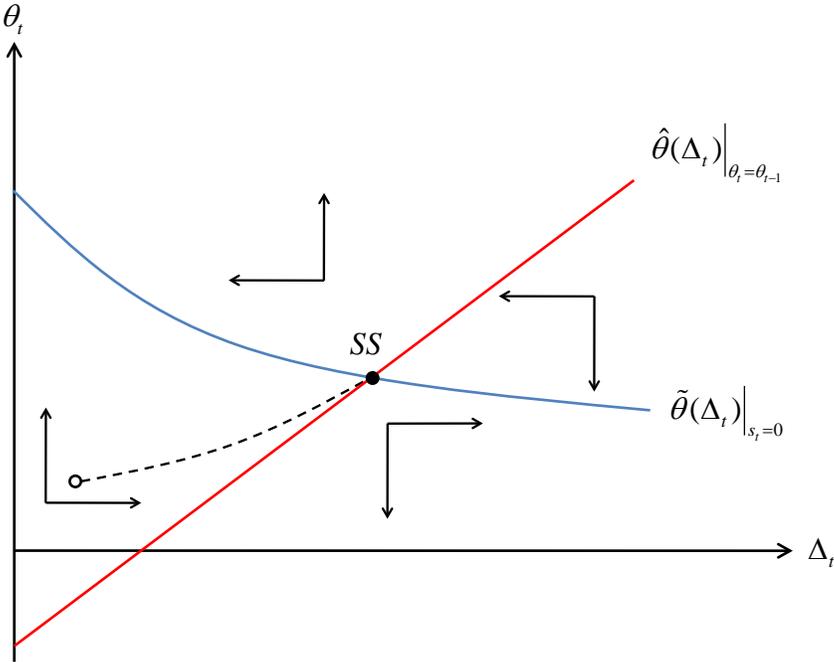


Figure 2: Change of steady state in response to underlying parameters

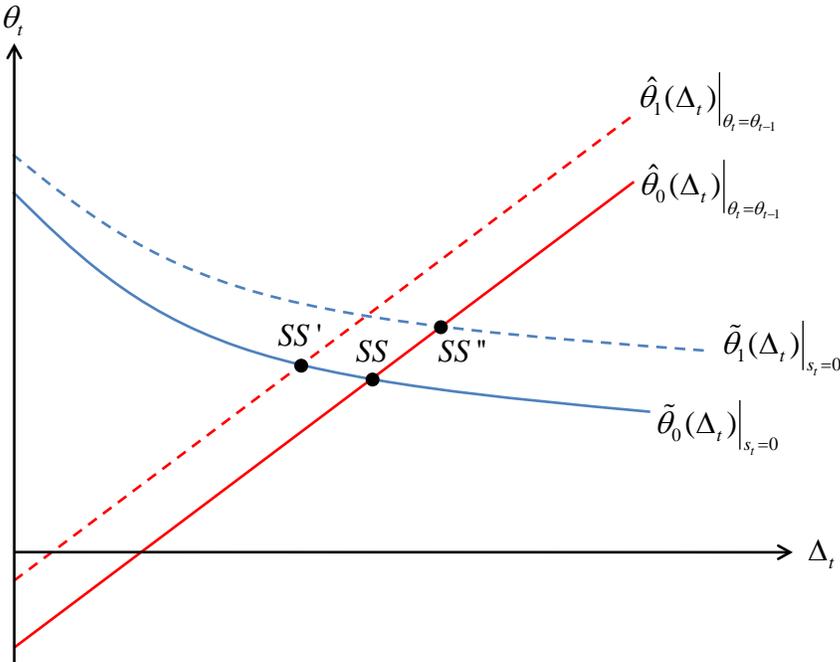


Figure 3: Time paths of the optimal solution

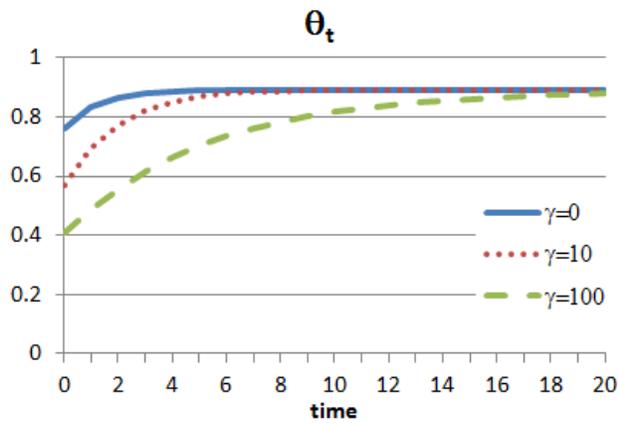


Fig. 3a: Cost limit

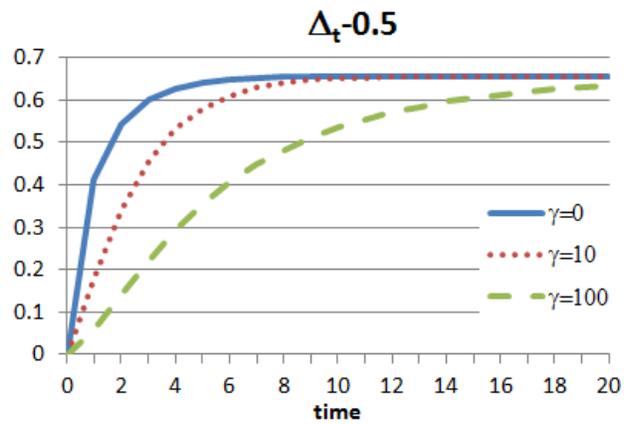


Fig. 3b: State of learning

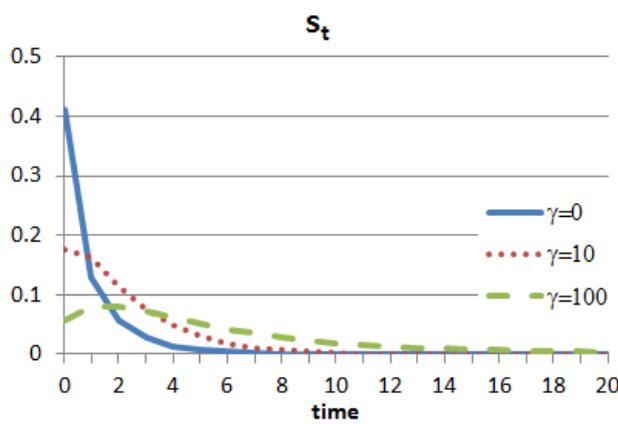


Fig. 3c: Rate of learning

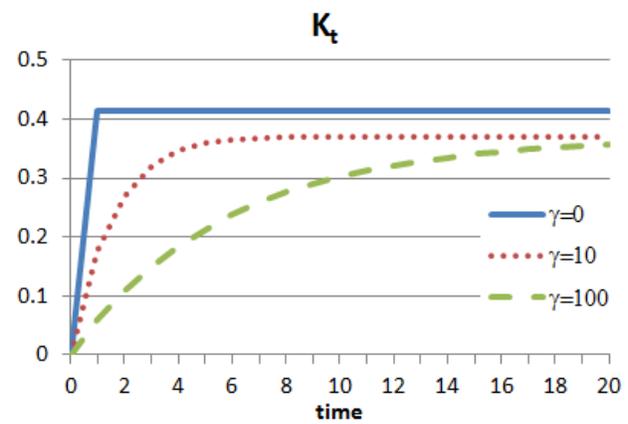


Fig. 3d: Total capacity

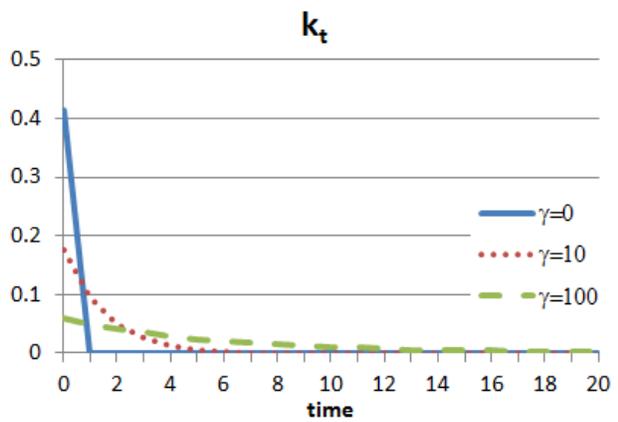


Fig. 3e: Rate of expansion

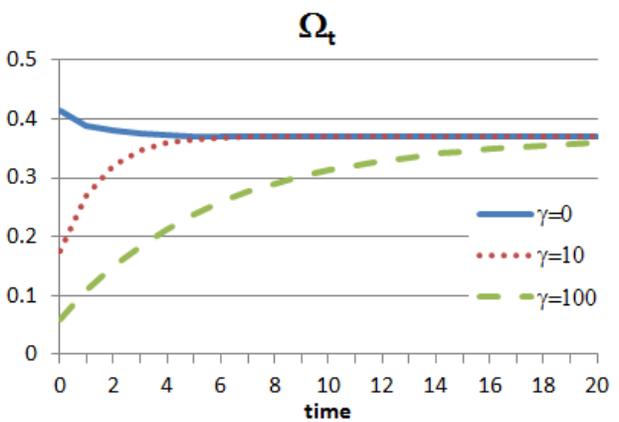


Fig. 3f: Usage rate

Appendix

Derivation of Eqs. (11) and (12)

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

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

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

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

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 (7) for $J_K(\Delta_{t+1}, K_{t+1})$:

$$\frac{\partial J(\Delta_{t+1}, K_{t+1})}{\partial K} = e^r (\lambda_t - Q'_t) \quad . \quad (\text{A.3})$$

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

Derivation of Eqs. (13) and (14)

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

$$\begin{aligned}
\theta_t|_{s_t=0; \lambda_t=\bar{\lambda}} &= D' \left[1 - \Omega_t|_{s_t=0} \right] - C' \left[\Omega_t|_{s_t=0} \right] - \bar{\lambda} \\
&= D' \left[1 - \int_{-\infty}^{\Delta_t} F(\theta_t - \Delta) dG(\Delta) \right] - C' \left[\int_{-\infty}^{\Delta_t} F(\theta_t - \Delta) dG(\Delta) \right] - \bar{\lambda}
\end{aligned} \tag{A.4}$$

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

$$\begin{aligned}
d\theta_t|_{s_t=0; \lambda_t=\bar{\lambda}} &= \frac{\partial \theta_t|_{s_t=0; \lambda_t=\bar{\lambda}}}{\partial \Delta_t} d\Delta_t + \frac{\partial \theta_t|_{s_t=0; \lambda_t=\bar{\lambda}}}{\partial \theta_t} d\theta_t|_{s_t=0; \lambda_t=\bar{\lambda}} \\
\frac{d\tilde{\theta}(\Delta_t)}{d\Delta_t} &\equiv \frac{d\theta_t|_{s_t=0; \lambda_t=\bar{\lambda}}}{d\Delta_t} = \frac{\partial \theta_t|_{s_t=0; \lambda_t=\bar{\lambda}} / \partial \Delta_t}{1 - \partial \theta_t|_{s_t=0; \lambda_t=\bar{\lambda}} / \partial \theta_t} \\
&= \frac{-(D_t'' + C_t'') F(\theta_t - \Delta_t) g(\Delta_t)}{1 + (D_t'' + C_t'') \int_{-\infty}^{\Delta_t} f(\theta_t - \Delta) dG(\Delta)} < 0
\end{aligned} \tag{13}$$

The relationship is negative due to our assumption that $D_t'' > 0$ and $C_t'' > 0$. Next, setting

$\theta_t = \theta_{t-1}$ in (11) and solving for θ_t leads to

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

Totally differentiating and rearranging gives:

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

Derivation of Eqs. (17) and (18)

Rearranging (16) leads to

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

Eq. (11) can be rearranged to

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

When substituting $\theta_t = \tau_t$, the RHS in (A.6) is identical to the LHS in (A.7) except for the interest rate. Combining (A.6) and (A.7) yields

$$\sigma_{t-1}e^{r_p} - \sigma_t = (e^{r_p} - e^r)(\Delta_t - \theta_{t-1}) \quad (\text{A.8})$$

Note that (A.1) when shifted by one period implies that $\Delta_t - \theta_{t-1} > 0$ as $J_\Delta(\Delta_t, K_t) < 0$ during the approach path, as the presence of more informed consumers strictly lowers the social costs if additional learning is still optimal. Note further that (11) implies for an interior steady state that

$$(e^r - 1 + F(\theta^{SS} - \Delta^{SS}))(\Delta^{SS} - \theta^{SS}) + \int_{-\infty}^{\theta^{SS} - \Delta^{SS}} \delta dF(\delta) = 0$$

which necessarily implies $\Delta^{SS} - \theta^{SS} > 0$ as the left hand-side is negative at $\Delta^{SS} - \theta^{SS} = 0$ and increasing in $\Delta^{SS} - \theta^{SS}$.¹⁹ Eq. (A.8) therefore implies that $\sigma_{t-1}e^{r_p} - \sigma_t$ has the same sign as

¹⁹ Intuitively, $J_\Delta(\Delta_\infty, K_\infty) < 0$ (and therefore $\Delta_\infty - \theta_\infty > 0$ by (A.1)) follows because there will always exist some consumers with $\delta < \theta_\infty - \Delta_\infty$ such that the social value of learning is strictly positive even in the steady state.

$e^{r_p} - e^r$ during the approach path and in the steady state. Eq. (A.8) further implies that the steady state subsidy σ^{SS} which induces type Δ^{SS} to learn is given by $\sigma^{SS} = (e^{r_p} - e^r)(\Delta^{SS} - \theta^{SS}) / (e^{r_p} - 1)$. For $r_p > r$ it follows that $\sigma^{SS} > 0$, and that the subsidy is positive along the whole approach path since $\sigma_{t-1}e^{r_p} > \sigma_t$. Note that after type Δ^{SS} has learned, the subsidy for first-time users could be reduced to zero, although this will have no effect on learning. For $r_p < r$, the subsidy is negative both in the steady state and along the approach path, i.e. it turns into a tax. For $r_p = r$ neither subsidy nor tax are required.

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

$$b_t \equiv \int_{-\infty}^{\theta_t - \Delta_t} (\Delta_t - \theta_t + \delta) d\tilde{F}(\delta) - \int_{-\infty}^{\theta_t - \Delta_t} (\Delta_t - \theta_t + \delta) dF(\delta), \quad (\text{A.9})$$

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

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

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

Proof of the Lemma

(i) The proof requires two steps (i.1) and (i.2).

(i.1) Assume that $\tau_t \leq \hat{\tau}(\Delta)$ and that tax rates $\tau_s \leq \hat{\tau}(\Delta)$ for all future periods $s > t$. Then learning at t generates weakly larger costs than never learning and is therefore not optimal. This is seen from the definition of $\hat{\tau}(\Delta)$ in (20) which implies:

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

(i.2). Part (i.1) immediately implies that learning in t can only be optimal if $\tau_t > \hat{\tau}(\Delta)$ or $\tau_s > \hat{\tau}(\Delta)$ for some periods $s > t$. To complete proving (i), we need to show that $\tau_t > \hat{\tau}(\Delta)$ is indeed necessary. To see this, assume that learning is optimal in t with $\tau_t \leq \hat{\tau}(\Delta)$ and define T as the first period in which $\tau_T > \tau_t$, i.e. $\tau_s \leq \tau_t$ for all $t \leq s < T$. Note that (i.1) implies that $T > t$ holds. We show that costs from learning in t are strictly higher than learning in T , implying that learning in t cannot be optimal:

$$\begin{aligned}
& \left[\Delta + \sum_{s=t+1}^T \left\{ \int_{-\infty}^{\tau_s - \Delta} (\Delta + \delta) dF(\delta) + (1 - F(\tau_s - \Delta)) \tau_s \right\} e^{-r_p(s-t)} \right] - \left[\tau_t + \sum_{s=t+1}^{T-1} \tau_s e^{-r_p(s-t)} + \Delta e^{-r_p(T-t)} \right] \\
&= \Delta - \tau_t + \sum_{s=t+1}^{T-1} \left\{ \int_{-\infty}^{\tau_s - \Delta} (\Delta + \delta - \tau_s) dF(\delta) \right\} e^{-r_p(s-t)} + \left(\int_{-\infty}^{\tau_T - \Delta} (\Delta + \delta - \tau_T) dF(\delta) - (\Delta - \tau_T) \right) e^{-r_p(T-t)} \\
&> \Delta - \tau_t + \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta - \tau_t) dF(\delta) \sum_{s=t+1}^{T-1} e^{-r_p(s-t)} + \left(\int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta - \tau_t) dF(\delta) - (\Delta - \tau_t) \right) e^{-r_p(T-t)} \\
&= (\Delta - \tau_t)(1 - e^{-r_p(T-t)}) + \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta - \tau_t) dF(\delta) \sum_{s=t+1}^T e^{-r_p(s-t)} \tag{A.11} \\
&\geq [\Delta - \tau_t] \left[1 - (e^{r_p} - 1) \sum_{s=t+1}^T e^{-r_p(s-t)} - e^{-r_p(T-t)} \right] \\
&= [\Delta - \tau_t] \left[1 - (e^{r_p} - 1) e^{-r_p} \frac{1 - e^{-r_p(T-t)}}{1 - e^{-r_p}} - e^{-r_p(T-t)} \right] \\
&= 0
\end{aligned}$$

Results (i.1) and (i.2) together prove claim (i) of Lemma.

(ii) Part (ii) of Lemma can now be easily seen: Assume that learning in t is optimal. Then, part (i) implies $\tau_t > \hat{\tau}(\Delta)$. To show that $\tau_t > \tau_{t-1}$, we compare the costs of learning in t with those of learning in in period $t-1$. Learning in t leads to lower costs than learning in $t-1$ if and only if

$$\begin{aligned}
& \Delta + e^{-r_p} \left\{ \int_{-\infty}^{\tau_t - \Delta} (\Delta + \delta) dF(\delta) + (1 - F(\tau_t - \Delta)) \tau_t \right\} > \tau_{t-1} + e^{-r_p} \Delta \\
&\Leftrightarrow e^{r_p} (\tau_t - \tau_{t-1}) > (e^{r_p} - 1)(\tau_t - \Delta) + \int_{-\infty}^{\tau_t - \Delta} (\tau_t - \Delta - \delta) dF(\delta)
\end{aligned} \tag{A.12}$$

Here, the right-hand side of the last inequality increases in τ_t and evaluates to zero in $\tau_t = \hat{\tau}(\Delta)$. Therefore it is greater than zero for $\tau_t > \hat{\tau}(\Delta)$ which implies $\tau_t > \tau_{t-1}$ and thereby completes the proof. ■

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:

$$\Delta_p + \frac{e^{-r_p}}{1 - e^{-r_p}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_p) - \Delta_p} (\Delta_p + \delta) dF(\delta) + (1 - F(\tilde{\theta}(\Delta_p) - \Delta_p)) \tilde{\theta}(\Delta_p) \right\} = \tau_T + \frac{e^{-r_p}}{1 - e^{-r_p}} \tilde{\theta}(\Delta_p) \quad (\text{A.13})$$

which can be solved for

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

The social costs from this path are given by:

$$\begin{aligned} \text{Cost}(\Delta_p) = & \left(\int_{-\infty}^{\Delta_p^{SS}} \int_{-\infty}^{\tau_T(\Delta_p) - \Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + \int_{\Delta_p^{SS}}^{\Delta_p} \Delta dG(\Delta) + D[1 - \Omega_T] + C[\Omega_T] + Q(\Omega_T - K_T) \right) \\ & + \frac{e^{-r}}{1 - e^{-r}} \left(\int_{-\infty}^{\Delta_p} \int_{-\infty}^{\tilde{\theta}(\Delta_p) - \Delta} (\Delta + \delta) dF(\delta) dG(\Delta) + D[1 - \Omega_p] + C[\Omega_p] \right) \end{aligned} \quad (\text{A.15})$$

$$\begin{aligned} \text{with } \Omega_T &= \int_{-\infty}^{\Delta_p^{SS}} F(\tau_T(\Delta_p) - \Delta) dG(\Delta) + (G(\Delta_p) - G(\Delta_p^{SS})) \\ \Omega_p &= \int_{-\infty}^{\Delta_p} F(\tilde{\theta}(\Delta_p) - \Delta) dG(\Delta) \end{aligned} \quad (\text{A.16})$$

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

$$\begin{aligned}
& \text{Cost}'(\Delta_p^{SS}) / g(\Delta_p^{SS}) \\
&= \Delta_p^{SS} - \underbrace{(D'(1-\Omega_T) - C'(\Omega_T))}_{=\tilde{\theta}(\Delta_p^{SS})} \\
&+ \frac{e^{-r}}{1-e^{-r}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_p^{SS})-\Delta_p^{SS}} (\Delta_p^{SS} + \delta) dF(\delta) - F(\tilde{\theta}(\Delta_p^{SS}) - \Delta_p^{SS}) \underbrace{(D'(1-\Omega_p) - C'(\Omega_p))}_{=\tilde{\theta}(\Delta_p^{SS})} \right\} \\
&= \Delta_p^{SS} - \tilde{\theta}(\Delta_p^{SS}) + \frac{e^{-r}}{1-e^{-r}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_p^{SS})-\Delta_p^{SS}} (\Delta_p^{SS} + \delta - \tilde{\theta}(\Delta_p^{SS})) dF(\delta) \right\} \\
&< \Delta_p^{SS} - \tilde{\theta}(\Delta_p^{SS}) + \frac{e^{-r_p}}{1-e^{-r_p}} \left\{ \int_{-\infty}^{\tilde{\theta}(\Delta_p^{SS})-\Delta_p^{SS}} (\Delta_p^{SS} + \delta - \tilde{\theta}(\Delta_p^{SS})) dF(\delta) \right\} \tag{A.17} \\
&= 0
\end{aligned}$$

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. ■