

Prices versus Quantities Reassessed^a

Larry Karp^b Christian Traeger^c

April 19, 2018

Abstract

“Prices versus quantities” (Weitzman 1974), a hugely influential paper, is widely cited (and taught) in current debates about the best policy to reduce greenhouse gas emissions. The paper’s criterion for ranking policies suggests that technological uncertainty favors taxes over cap and trade. Weitzman models a flow pollutant, but greenhouse gases are persistent. Stock pollutants require a fundamental change in the ranking criterion. Innovations’ persistence and their gradual diffusion both favor the use of cap and trade. Numerical results show that the case for cap and trade as a means of reducing greenhouse gas emissions is stronger than widely believed.

JEL Codes: Q00, Q50, H20, D80

Keywords: policy instruments, pollution, climate change, taxes, quantities, regulation, uncertainty, cap and trade, technology

^aWe are grateful for comments by Hiroaki Sakamoto, Reyer Gerlagh, Inge van Den Bijgaart, Rolf Golombek, and seminar participants at the EAERE 2017, SITE 2017, and CES-ifo area network conference 2017. CT gratefully acknowledges financial support by the Norwegian Research Council.

^bDepartment of Agricultural and Resource Economics, UC Berkeley

^cDepartment of Economics, University of Oslo; ifo Institute for Economics Research; Frisch Centre.

1 Introduction

Weitzman's (1974) "Prices versus quantities" is among the most widely taught papers in environmental economics. It is elegant, simple, and makes an important point about ranking the two most popular market-based policy instruments to control pollutants. The paper has been widely cited in debates over climate policy (total citations > 3500). In teaching students about climate change and cap-and-trade systems, we use Weitzman's (1974) reasoning to explain why society would probably be better off using taxes.

In Weitzman's setting, regulators are uncertain about firms' abatement costs. He shows that this uncertainty creates a lower deadweight loss under a tax than under cap and trade (a quota) if and only if the slope of the marginal abatement cost curve is steeper than the slope of the marginal damage curve. His paper describes a flow pollutant, making the analysis simple and crisp. A flow pollutant affects society only in the period when it is released. However, most regulated pollutants have some persistence or cumulative impact. For example, carbon dioxide emissions have an effective half-life well above a century. Weitzman's logic underpins much of the discussion about instrument choice in climate policy. We show that this logic fails for the stock pollutants causing climate change.

The correction needed to rank taxes and quotas for stock pollutants is almost as simple as Weitzman's criterion for a flow pollutant. Section 2 explains the basic idea using graphs. Section 3 presents the formal results using a dynamic programming model with serially correlated technological shocks and gradual diffusion of new technologies, under asymmetric information. Section 4 calibrates the model based on widely-used estimates of abatement costs and climate-related damages. It shows that the advantage of taxes over quotas in controlling greenhouse gas emissions is much weaker than widely believed.

We now review the intuition for Weitzman's ranking and explain the correction needed for stock pollutants. Under a tax, technological uncertainty creates uncertainty about emissions, and consequently about damages. Under a binding quota, technological uncertainty does not alter emissions, but it creates cost uncertainty for firms. If the (aggregate) marginal abatement cost curve is steeper than the marginal damage curve, then uncertainty about abatement costs harms society more than uncertainty about damages. Hence, in the static (flow pollutant) setting taxes are preferred if and only if the marginal abatement cost curve is steeper than the marginal damage curve.

In a dynamic setting with stock pollution, in each period we have to compare current abatement costs with the discounted stream of damages occurring over an extended time horizon. Hereafter, in discussing stock pollutants,

we refer to the discounted stream of marginal damages simply as “marginal damages”. In the climate context, the relevant marginal damage function is referred to as the social cost of carbon. Applying the logic described above to the climate context, one might be tempted to rank taxes and quotas by comparing the slopes of the social cost of carbon and the marginal abatement cost curve. Indeed, the literature makes this leap (e.g. Nordhaus 2008¹, Wood & Jotzo 2011² Weitzman 2018³).⁴

This leap is generally incorrect. In the dynamic setting the abatement technology is a stochastic process. Every period gives rise to an innovation unknown to the policy maker at the onset of the regulation period. Because technology is highly persistent, the technological innovation affects not only the abatement cost in the present period, but also the abatement costs in future periods. Consequently, the innovation also changes future optimal (or business as usual) emissions, thus changing the future baseline concentration of the pollutant. With convex damages, the marginal damage caused by additional emissions today depends on the pollutant’s baseline concentration.

As a consequence, the technology shock not only shifts the marginal abatement cost curve, but it simultaneously shifts the marginal damage curve (= social cost of carbon). The policy ranking depends as much on this shift as it does on relative slopes. In an extreme case, this additional shift implies that a cap and trade system is first best (not just better than the tax) even if the marginal damage curve is much flatter than the abatement cost curve.

Technological innovations typically diffuse slowly through the economy, rather than being adapted instantly (Allan et al. 2014). Gradual diffusion creates a gap between the feasible and the implemented technology level; current adoption signals greater future adoption. With gradual diffusion, a technology shock might have little effect on current abatement costs, but the future adoption can lead to a large shift in the social cost of carbon. Thus, gradual diffusion of technology also favors quotas over taxes.

¹“Uncertainty pervades climate-change science, economics, and policy. One key difference between price and quantity instruments is how well each adapts to deep uncertainty. A major result from environmental economics is that the relative efficiency of price and quantity regulation depends upon the nature – and more precisely the degree of nonlinearity of costs and benefits (see Weitzman 1974).”

²“It is generally thought that this is the case with climate change for the comparison between price and quantity instruments.”

³“For example in the case of CO₂, since the marginal benefit curve within a regulatory period is very flat [...] the theory strongly advises a fixed price as the optimal regulatory instrument.”

⁴Other recent reviews of taxes and quotas in the context of climate policy, include Hepburn (2006), Aldy et al. (2010), Goulder & Schein (2013), and Newbery (2018).

Hoel & Karp (2002) analyze prices versus quantities for a stock pollutant in a setting where current innovations have no impact on future technology, thereby ruling out our results.⁵ Newell & Pizer (2003) rank the two policies when cost shocks are serially correlated in an open-loop setting, where the regulator chooses current and future policy levels at the initial time. They show that positively correlated cost shocks increase stock volatility under taxes, favoring quotas, although not by enough to tip the balance.⁶ The intuition they provide depends crucially on the open loop assumption that future policy makers do not respond when the realized technology and emission levels drift away from anticipated levels. In contrast, the effect we describe precisely depends on the fact that policy makers will update the emission targets after observing past technological innovations.

Karp & Zhang (2005) compare the policy ranking across the open loop and feedback settings, with correlated shocks. They observe that the ranking of taxes versus quantities changes, but they do not recognize the different mechanism in the feedback setting and its ability to make quotas dominate taxes in the climate context. Our contributions are (i) to develop a simple and intuitive ranking criterion that (ii) explains why Weitzman's reasoning does not carry over to stock pollutants, and to show that (iii) the case for taxes is much weaker in climate change than previously thought, and that additionally (iv) under slow technology diffusion quotas can, in special cases, implement the first best allocation under uncertainty (even when the marginal abatement cost curve is steeper than marginal damages).

2 One-period graphical analysis

Weitzman's static model for a flow pollutant produces a simple criterion for ranking a tax and quota. A variation of this one-period model reveals a fun-

⁵Karp & Zhang (2012) study a more general model that includes endogenous investment in abatement capital.

⁶Under taxes in the open loop setting, positively serially correlated shocks translate into positively serially correlated levels of emissions. These raise the volatility of the pollution stock and increase the deadweight loss arising from stock uncertainty. In contrast, under quotas in the open loop setting, the stock trajectory is deterministic. Serial correlation does not have a similar impact on the deadweight loss arising from abatement cost uncertainty, because abatement costs depend only on a period's shock realization, not on its history. This intuition breaks down under feedback policies, where the stock trajectory is stochastic under both taxes and quotas. By conditioning future policies on historic shock realizations, policy makers eliminate the cumulative deviation between the realized and the optimal stock levels.

damental difference between the settings where damages depend on the flow of pollution or the stock of pollution. The criterion for ranking policies in the stock-related case is only slightly more complicated than in the flow-related case, and it closely relates to the formula we develop for the dynamic model.

2.1 Review of standard model

In the classic prices versus quantities setting, marginal damages increase linearly in emissions: $MD = a + bE$. The slope parameter b characterizes the convexity of damages. Similarly, the classical setting assumes that marginal benefits from emissions are linear. An optimizing firm emits to the point where the marginal benefits of emissions equal the marginal abatement costs. We write these marginal costs as a function of emissions (instead of abatement): $MAC = \theta - fE$. The slope parameter f captures the concavity of the benefits from emitting or, equivalently, the convexity of the abatement cost.⁷ The upper left panel in Figure 1 depicts the MD curve as the increasing solid line and shows the expected abatement cost curve as the decreasing solid line.


The parameter θ is private information, known to the firm but not to the policy maker. The planner knows only the expected value of θ . A risk neutral planner sets $\mathbb{E}(MAC) = MD$, equating the marginal damage curve and the expectation of the marginal abatement cost curve.⁸ With taxes, the policy fixes the emissions price at the green (horizontal) line in Figure 1. In a quantity setting, the policy caps the emissions at the red (vertical) line.

Figure 1 shows the realized marginal abatement cost curve as the dashed line. The top left panel shows the tax and the quota equilibria for a flow pollutant, both of which differ from the social optimum. The figure identifies the deadweight loss under the tax as a light green triangle and the deadweight loss under a quota as a red triangle. The (green) deadweight loss under the tax

⁷Let the absolute benefits of emissions be $B(E) = \theta E - \frac{f}{2} E^2$. Abatement is the difference between business as usual and actual emissions: $A = E^{BAU} - E$. Business as usual emission are industry's optimal emissions in the absence of policy. Firms' first order condition for unregulated emission optimization yields $E^{BAU} = \frac{\theta}{f}$. Thus, the absolute abatement costs are $AC(A) = B(E^{BAU}) - B(E) = \theta E^{BAU} - \frac{f}{2} E^{BAU^2} - \theta E + \frac{f}{2} E^2 = \frac{1}{2} \frac{\theta^2}{f} - \theta E + \frac{f}{2} E^2$ resulting in the marginal abatement costs $MAC(A) = (-\theta + fE) \frac{dE}{dA} = \theta - fE$. Thus, f indeed describes both the concavity of emission benefits and the convexity of abatement costs.

⁸The common assumption that the intercept but not the slope is private information is key to the simplicity of both Weitzman's and our result (Perino & Requate (2012)). Hoel & Karp (2001) rank the two policies in a model with stock pollutants, when a serially uncorrelated shock affects the slope. The resulting criterion for policy ranking is not closely related to the criterion where the shock affects the intercept of marginal cost.

Flow pollution



Stock pollution




Figure 1: Illustration of Weitzman (1974) insights for a flow pollutant (top panels) and a quasi-static illustration of the changes for a stock pollutant (bottom panels). The light green (left) triangle characterizes the deadweight loss under a tax, whereas the red (right) triangle characterizes the deadweight loss under a quota. The black solid lines represent expectations, and their dashed counterparts represent realizations. The panels on the right add labels of relevant distances and slopes for our graph-based quantitative illustration of taxes versus quotas.

is smaller than the (red) deadweight emerging under a quota. In this figure, taxes dominate quotas because the *MAC* curve is steeper than the *MD* curve.

2.2 Modification for a stock pollutant

Weitzman's (1974) setting assumes that a change in θ does not shift the MD curve. A footnote in his paper, elaborated by Stavins (1996), points out that a correlation between θ and the MD curve complicates the policy ranking. Stavins describes situations where an underlying factor simultaneously affects marginal abatement costs and marginal damages. For example, a sunny day increases ultraviolet radiation, increasing ozone production, raising ozone abatement costs. If the sunny day causes people to spend more hours outdoors, marginal damages from ozone (respiratory stress) also increase.

In a dynamic setting, stock pollutants and abatement technologies evolve over time. The policy maker has to regulate pollution without knowing the current or future abatement costs. Technology shocks are persistent and an innovation simultaneously lowers current and future abatement costs. A reduction in the future costs changes future emissions, altering future levels of the stock pollutant. The change in future stocks affects the discounted stream of damages resulting from emitting an additional unit of the pollutant today. Thus, the innovation that shifts the marginal benefits curve *also shifts the marginal damage curve*. Here we need not search for a distinct factor (e.g. the sunny day in Stavins's example) that affects both abatement costs and damages. The uncertainty about technology at the heart of Weitzman's (1974) framework necessarily affects marginal damages as well.

The lower left panel in Figure 1 illustrates the consequences of the above insight. The MD curve now represents the discounted stream of future damages arising from the flows' impact on the pollution stock: the social cost of carbon in the climate context. To use graphical analysis, we take this marginal damage function as exogenous. In the genuinely dynamic setting discussed in Section 3, the marginal damage function is endogenous.

We again assume that an innovation lowers, by Δ , the marginal abatement cost from the solid line to the dashed line. Because technology is persistent, the reduction in marginal abatement costs makes future emission reductions cheaper, and reduces future emissions (both optimal and business as usual). The resulting reduction in the future pollution stock lowers the marginal damage from releasing an additional unit of the pollutant today. As a consequence, the MD curve also shifts down. The parameter φ relates the shifts of the MD and the MAC curves. In the climate change setting, φ equals the derivative of the social cost of carbon with respect to the technology shock.

The lower panels in Figure 1 show that the downward shift of the MD curve increases the deadweight loss of the tax and reduces the deadweight loss of the quota. The ranking of taxes versus quantities now depends on both the

relative slopes of the two curves and the shifter φ .

2.3 Policy ranking with a stock pollutant

This subsection offers a graphical perspective on some results derived in the dynamic model (Section 3). The right panels in Figure 1 enrich the graphs of the left panels by adding labels to the slopes and some of the distances in the graphs.

We use the upper right panel of Figure 1 to establish **Weitzman's result** that a tax dominates a quota for a flow pollutant if and only if the *MAC* curve is steeper than the *MD* curve. The green and the red triangles representing the deadweight losses are similar, i.e., they have the same angles. As a consequence, the deadweight loss of the tax (light green triangle) is smaller than the deadweight loss of the quota (red triangle) if and only if $s < k$. The absolute value of the slope of the *MAC* curve is $m^{MAC} \equiv \frac{k}{h}$ and the slope of the *MD* curve is $m^{MD} \equiv \frac{s}{h}$. Thus, $\frac{m^{MAC}}{m^{MD}} = \frac{k}{s}$ and taxes dominate quotas ($\frac{k}{s} > 1$) if and only if $m^{MAC} > m^{MD}$, confirming Weitzman's result.

The lower right panel in Figure 1 adds labels to the one-period illustration for the **stock pollutant**. Now, in addition to the slopes, the responsiveness φ of the *MD* curve to a given shift of the *MAC* curve matters. We use three geometrical relations from the graph. First, we relate the deadweight loss under the quota to the relative shift φ . Using the relation $\frac{d}{d+k} = \varphi$, we have

$$d = \frac{\varphi k}{1 - \varphi} \quad \text{or} \quad \frac{d}{k} = \frac{\varphi}{1 - \varphi}. \quad (1)$$

Once again, the light green and the red triangles representing the deadweight loss in the two settings are similar (same angles), and we compare them based on their sides s and k . By the definition of the slope, $h m^{MAC} = k + d \Rightarrow \frac{h m^{MAC}}{k} = 1 + \frac{d}{k}$. Using equation (1) to replace the fraction $\frac{d}{k}$ delivers

$$\frac{h m^{MAC}}{k} = \frac{1}{1 - \varphi}. \quad (2)$$

Similarly, we observe that $h m^{MD} = s - d \Rightarrow \frac{h m^{MD}}{s} = 1 - \frac{d}{s}$. Using equation (1) to replace d , we obtain

$$\frac{h m^{MD}}{s} = 1 - \frac{\varphi k}{1 - \varphi s}. \quad (3)$$

Dividing equation (3) by equation (2) and solving for $\frac{s}{k}$ delivers

$$\frac{s}{k} = \frac{1}{1 - \varphi} \left(\frac{m^{MD}}{m^{MAC}} + \varphi \right).$$

Taxes dominate quotas if and only if the deadweight loss of the tax is smaller than the deadweight loss of a quota, i.e., $\frac{s}{k} < 1$, which is equivalent to

$$\frac{m^{MD}}{m^{MAC}} < 1 - 2\varphi \quad \Leftrightarrow \quad \frac{m^{MAC}}{m^{MD}} > \frac{1}{1 - 2\varphi} \gg 1. \quad (4)$$

If the marginal damage curve is somewhat responsive to the technology shock, then the *MAC* curve has to be much larger than the *MD* curve for taxes to be the preferred instrument. The next section builds an explicit model of stock pollution, verifying and extending equation (4). It finds a closely related formula and expresses the slope of marginal damages and the shifter φ in terms of fundamentals.

3 The dynamic model

Many types of shocks or trends can affect the unregulated level of emissions, thus affecting the cost of abatement. The policy ranking depends on asymmetric information between the regulator and firms, i.e., on private information. Publicly observed shocks that are uncorrelated with private information do not affect the policy ranking, only policy levels. We ignore such shocks, but Appendix A discusses their inclusion.⁹ After describing the model we discuss the policy ranking, extending the insights from the previous section to a dynamic setting.

3.1 Description of the model

The model consists of two state variables, a pollution stock and a technology level. We express the pollution stock, S_t , as the deviation from the “no-harm”

⁹Instead of choosing a policy level, a tax or a quota, at the beginning of each period, the regulator could announce a policy rule conditioned on the realizations of publicly observed random variables occurring during the regulatory period. Here, the policy level, but not the policy rule, adjusts within a regulatory period. Most importantly, policy levels should explicitly account for the emissions impact of business cycles by explicitly conditioning on macroeconomic observables.

level; for climate change S_t corresponds to the deviation from the pre-industrial carbon concentration. The equation of motion for the stock is

$$S_{t+1} = \delta S_t + E_t,$$

where E_t is the annual emission flow and $\delta \in (0, 1)$ captures the persistence of harmful pollution (decay factor $1 - \delta$). The second state variable, θ_t , governs firms' feasible abatement technology. Abatement costs also depend on h_t , an exogenous trend; θ_t represents the stochastic deviation from this trend. This deviation is a highly persistent stochastic process under iid shocks $\varepsilon_t \sim iid(0, \sigma^2)$. The equation of motion for θ is

$$\theta_t = \rho\theta_{t-1} + \varepsilon_t, \text{ with } \rho > 0 \text{ and } \mathbb{E}_t(\varepsilon_t) = 0.$$

The policy maker knows θ_{t-1} but not ε_t when she chooses policy at t ; firms know both θ_{t-1} and ε_t at time t . This asymmetry provides the dynamic analogue of Weitzman's (1974) asymmetric information. The abatement cost parameter θ_t characterizes the deviation from the expected technology level.

The pollution stock causes annual damages of $\frac{b}{2}S_t^2$. The exogenous parameter b equals the slope of the marginal *flow* damage curve. The (discounted stream of) the marginal damage from releasing another unit of emissions and its dependence on technology shocks is endogenously determined by the model, not an exogenous input as in the preceding section. The policy maker balances the persistent costs from pollution with the transitory benefits from emitting.

Firms' response to a technology shock ε_t can take some time. We use the parameter α to distinguish between the currently feasible level of technology, θ_t , and the level adopted by firms, $\hat{\theta}_t$:

$$\hat{\theta}_t \equiv \rho\theta_{t-1} + \alpha\varepsilon_t. \tag{5}$$

Firms immediately adopt only the share $\alpha \in (0, 1]$ of the current shock; they adopt the remaining fraction in the next period. This assumption captures, without the need of an additional state variable, the well-established fact that technology diffuses gradually: an uptake in adoption today can signal additional adoption tomorrow. For $\alpha < 1$, the correlation between $\hat{\theta}_t$ and $\hat{\theta}_{t+j}$ is greater, and falls more slowly with j , relative to the model with $\alpha = 1$.¹⁰

The emission benefits (or abatement costs) depend on $\hat{\theta}_t$: $(h_t + \hat{\theta}_t)E_t - \frac{f}{2}E_t^2$, where f is the slope of the marginal abatement cost curve. A higher value of

¹⁰The correlation between $\hat{\theta}_t$ and $\hat{\theta}_{t+j}$ is $\rho^j \left(\rho^2 \frac{\rho^{2j} - 1}{\rho^2 - 1} + \alpha^2 \right)^{-0.5}$, see Appendix B.3.

$\hat{\theta}_t$ corresponds to a larger marginal benefit from emitting: a larger marginal abatement cost. Hereafter we reserve the term “innovation” to mean a negative value of the shock ε : an innovation lowers the marginal benefit of emissions, thus lowering abatement costs.

We use superscripts Q and T for the quota and tax policy scenarios. Under a binding **quota**, the regulator chooses the actual emissions level E_t^Q and has the *expected flow net benefit* (using $\mathbb{E}_t \alpha \varepsilon_t = 0$)

$$(h_t + \rho\theta_{t-1}) E_t^Q - \frac{f}{2} (E_t^Q)^2 - \frac{b}{2} S_t^2.$$

Under a **tax** τ_t the *firm's payoff* is $(h_t + \hat{\theta}_t)E - \frac{f}{2}E_t^2 - \tau_t E_t$, implying the first order condition $h_t + \hat{\theta}_t - fE_t = \tau_t$. Using the definition (5), the firm's decision rule is

$$E_t^T = e_t^T + \alpha \frac{\varepsilon_t}{f} \quad \text{with} \quad e_t^T \equiv \frac{h_t + \rho\theta_{t-1} - \tau_t}{f} \quad \left(= \mathbb{E} E_t^T \right).$$

It is convenient to model the tax-setting regulator as choosing *expected emissions* e_t^T , which is equivalent to setting the tax $\tau_t = h_t + \rho\theta_{t-1} - f e_t^T$. The tax payment is a pure transfer and does not enter the regulator's payoff function. The tax-setting regulator's *expected flow net benefit* from emissions is

$$(h_t + \rho\theta_{t-1}) e_t^T - \frac{f}{2} (e_t^T)^2 + \frac{\alpha^2}{2f} \sigma^2 - \frac{b}{2} S_t^2.$$

For both problems, the regulator wants to maximize the expectation of the present discounted stream of net benefit flows, defined as the benefit of emissions minus the stock-related damage. The discount factor is β .

3.2 Policy ranking

The marginal pollution damage, the Social Cost of Carbon (*SCC*) in the climate setting, is

$$SCC_t = \chi_t + \lambda S_t + \mu \theta_{t-1}.$$

Appendix A provides formulae for λ , the derivative of the *SCC* with respect to the carbon stock, and μ , the derivative of the *SCC* with respect to the technology parameter, and shows that both are positive constants. An increase in the stock of carbon increases the *SCC*, and (because $\rho > 0$) a higher cost of abatement shifts up the graph of the *SCC* as a function of carbon. The time

dependence of χ_t reflects the *SCC*'s response to the technology trend h_t . The functions χ_t , λ and μ are the same under the optimal tax, the optimal quota, and in the full information (first best) setting where the planner observes ε_t . The levels of emissions are the same in the three settings if and only if the shock equals its expected value, $\varepsilon_t = 0$.

We denote by $r \equiv \frac{b}{f}$ the ratio of the slopes of the marginal flow damage and the marginal abatement cost. This slope describes the relative convexity of the damage and the abatement cost functions. For a flow pollutant, taxes dominate quotas if and only if $r < 1$. In the case of carbon dioxide, r is tiny, about 8×10^{-4} for our baseline calibration (Section 4.1). For the case of a stock pollutant, the intertemporally aggregated marginal damages, the *SCC*, replace the flow marginal damages. Accordingly, we define the ratio $R \equiv \frac{\lambda}{f}$, which relates the convexity of stock damages to that of abatement costs.¹¹ Lemma 1 gives the relation between these two slopes.

Lemma 1 *Under both taxes and quotas, the slope of the SCC with respect to the stock of carbon, relative to the slope of marginal abatement cost is*

$$R \equiv \frac{\lambda}{f} = \frac{1}{2\beta} \left(- (1 - \beta\delta^2) + \beta r + \sqrt{(1 - \beta\delta^2 - \beta r)^2 + 4\beta r} \right). \quad (6)$$

Unsurprisingly, the discount factor β and the persistence of the pollutant δ play the major role in relating the flow ratio r and the stock ratio R . Figure 2 graphs $\frac{R}{r}$ as functions of the flow pollution ratio r for our baseline calibration of β and δ (solid) and for an alternative with higher discount factor and higher persistence (dashed). In our base calibration, R is 23 times as large as r . This factor means that aggregate damages are *more convex* than flow damages: the *SCC* is much steeper in emissions than is the flow marginal damage curve.

The following proposition provides two equivalent characterizations of the criterion ranking taxes and quotas for a stock pollutant.

Proposition 1 *Taxes dominate quotas if and only if*

$$R < \frac{1}{\beta} - \frac{2\mu}{\alpha} \quad \Leftrightarrow \quad R < R^{crit} \equiv -\frac{1}{2}\kappa_1 + \frac{1}{2}\sqrt{\kappa_1^2 + 4\kappa_0} \quad (7)$$

$$\text{with } \kappa_1 \equiv \frac{\delta\rho(2-\alpha)}{\alpha} \text{ and } \kappa_0 \equiv \frac{1-\beta\delta\rho}{\beta^2}.$$

¹¹ Appendix A uses a parameter ϕ to denote the number of years in a period, producing a simple means of accommodating a flexible time step. There we have the definition $R \equiv \frac{\lambda}{f}\phi$. Here, λ still captures the value of a unit increase in atmospheric carbon, but f captures the benefits from an annual emission flow over the course of the period. Thus the derivative of abatement cost with respect to an absolute change of emissions is $\frac{f}{\phi}$.




Figure 2: The solid graph shows the ratio $\frac{R}{r}$ for our baseline values of β and δ with an annual time step. The dashed graph increases the half life of the stock by 50% and lowers the discount rate from 2% to 1% per year. “ r^{base} ” marks the slope ratio in our baseline calibration.

For flow pollutants taxes dominate quotas if and only if $r < 1$. The first condition in Proposition 1 shows that (i) the relevant slope criterion for a stock pollutant is R instead of r , (ii) a higher value of μ favors quotas, and (iii) slow technology adoption favors quotas.¹² The shadow value μ captures the responsiveness of the social cost of marginal emissions (SCC) to the technology level. This shadow value is endogenous to the model. The right hand side of the equivalence (7) expresses the criterion in terms of the fundamental model parameters. We note that the ratio R and the critical level R^{crit} respond differently to parameters: R depends on all parameters except α , whereas R^{crit} depends on all parameters except r .

Figures 3 graphs R^{crit} as a function of the “joint persistence” $\delta\rho$ of the stock pollutant and technology for three values of α , assuming an annual time step and the discount factor $\beta = 0.98$. For the case of no persistence of pollution or technology $\delta\rho$ is close to zero and the left panel shows that the critical value is close to unity, as in the static criterion. However, for climate change δ is close to 1; and with persistent technology so is ρ . For $\delta\rho \approx 1$ the right panel of Figure 3 shows that the critical value remains bounded away from 0. In the climate change context, quotas might dominate taxes not only when r is tiny,

¹²Equation (12) in the appendix provides the formula for μ in terms of the the model’s fundamentals. Importantly, μ is independent of α .




Figure 3: R^{crit} as a function of $\delta\rho$ for $\alpha = 1$ (red dash), $\alpha = 0.75$ (solid), and $\alpha = 0.5$ (green dash), with an annual time step. Quotas dominate taxes for $R > R^{crit}$. The upper right graph zooms in on high values of $\delta\rho$ as they prevail in the climate context, the “o” marks our baseline calibration.

but even if R is close to 0. Section 4 further explores this possibility.

We provide intuition for our results building on the case of a *flow pollutant*. There, a technology innovation (a negative value of ε) lowers both the socially optimal emission level and marginal abatement cost. Under taxes, firms face constant abatement prices and the emission quantity overreacts compared to the socially optimal response. This quantity fluctuation’s impact on expected damages is the dominating contribution to the deadweight loss under a tax. By Jensen’s inequality, the convexity of the damage function determines its magnitude. Under quotas, emissions are constant but the firms’ marginal abatement costs overreact compared to the social optimum. This exaggerated abatement cost fluctuation is the dominating contribution to the deadweight loss under quotas, and by Jensen’s inequality the convexity of the abatement cost function determines its magnitude. If abatement costs are more convex than damages, the deadweight loss is larger under a quotas. Lemma 1 shows that *stock pollution* increases the relative convexity of the damage function (R instead of r), thereby weakening the case for taxes.

Proposition 1 states that, for a stock pollutant, a greater sensitivity of the SCC to technology (higher μ) and a slower technology diffusion (smaller α) further *strengthen the case for quotas*. We start by providing the intuition for the case of immediate technology diffusion ($\alpha = 1$). As discussed in the preceding paragraph, the dominating contribution to the deadweight loss under quantity

regulation of flow pollutants is the overreaction of the equilibrium abatement price relative to the socially optimal response. For a stock pollutant, technological innovation today implies lower future emissions, resulting in a lower future pollution stock.¹³ Consequently, a technological innovation *reduces the marginal damages* resulting from an additional emission unit today (assuming convex damages). This reduction in marginal damages amplifies the socially optimal price fluctuation resulting from the innovation's reduction of marginal abatement costs. Thus, the socially optimal price fluctuation is larger in the stock pollution setting than in the flow pollution setting: the innovation affects marginal costs and marginal damages in a perfectly correlated way. A part of what would be an "overreaction" of abatement prices under quantity regulation of a flow pollutant becomes a socially optimal variation under a stock pollutant.

Proposition 1 shows that a higher (endogenous) value of μ favors quotas. The shadow value μ of the interaction term $\theta_t S_t$ measures the responsiveness of the social cost of marginal emissions to the technology level. If the socially optimal abatement cost responds more sharply to innovation (μ large), then the socially optimal response approaches the "overreaction" of abatement costs under a quantity regulation, reducing the deadweight loss of a quota. The graphical analysis in Section 2 reflects this intuition. When the technological innovation shifts the marginal damage curve for a stock pollutant (lower panels of Figure 1), it amplifies the optimal price fluctuations in response to the innovation, relative to the case of the flow pollutant (upper panels of Figure 1). Indeed, for $\alpha = 1$, the left side of the policy-ranking equivalence (7) (dynamic model) reproduces the left side of the graph-based equivalence (4) that we derived in the quasi-static setting. The dynamic model introduces the additional discount factor only because we assume that today's emissions contribute to tomorrow's stock and damages, whereas the quasi-static analog treated the damage as instantaneous.

The main difference between the quasi-static stock pollution extension and the dynamic model is that both R and μ are endogenous in equation (7), whereas the quasi-static model simply assumed some slope ratio of marginal damages over marginal abatement costs and merely argued for the existence of some shift, φ , of the marginal damage curve. In addition, the quasi-static model cannot capture the fact that technology diffusion takes more than one period ($\alpha < 1$). Before continuing the discussion of technology diffusion and

¹³In line with the empirical findings for most sectors, our functional forms imply that there is no rebound effect strong enough to increase aggregate emissions in response to an emissions-saving innovation




Figure 4: Quotas dominate taxes for $\alpha < \alpha^{crit}$, shown as the solid graph. The quota is first best for $\alpha = \alpha^*$, shown as the red dashed graph. The graphs use an annual time step and baseline values $\beta = 0.98$, $\delta = 0.997$, and $\rho = 0.99$. “ R^{base} ” marks the slope ratio in our baseline calibration.

the underlying intuition we pose one more question. Can the “overreaction” of marginal abatement costs from the flow pollution perspective become a socially optimal fluctuation for a stock pollutant?

Proposition 2 *Assume that $b, f, \beta, \rho, \delta > 0$ and that $\beta\delta\rho < 1$.*

- (i) *There exists $\alpha^* \in (0, 1)$ such that the quota is first best.*
- (ii) *A reduction in α favors quotas, and there exists $\alpha^{crit} \in (\alpha^*, 1)$ such that quotas dominate taxes for all $\alpha < \alpha^{crit}$.*

The proposition shows that for any model calibration with convex damages and abatement costs there exists a technology adoption rate α for which quotas dominate taxes. For sufficiently slow technology diffusion, quotas are not only preferred to taxes, but the cap and trade system achieves the first best emission allocation, even if the slope of the marginal damages curve is arbitrarily small (but positive) and the slope of the marginal cost curve is arbitrarily large. The proposition also implies that this situation can arise only under partial technology diffusion ($\alpha < 1$). Figure 4 graphs α^* , where quotas are first best, and α^{crit} , below which quotas dominate taxes, as a function of R , the ratio of stock damage convexity to abatement cost convexity.

To understand the role of technology diffusion, note that under partial diffusion today's technology shock provides information not only about today's technology adoption but also about subsequent adoption. Footnote 10 and the discussion below equation (5) note that partial diffusion increases the correlation between the current and future adopted technology levels. As a result, a given level of adoption today signals even more future adoption. In particular, the socially optimal marginal abatement cost responds to innovation, anticipating both present and future adoption. In contrast, the marginal abatement cost under quantity regulation responds only to the presently adopted part of the innovation. As a consequence, partial diffusion increases fluctuations of the socially optimal price of emissions relative to the fluctuations arising under quantity regulation. Given that quantity regulation generally suffers from an overreaction of the emissions price, partial diffusion reduces the welfare loss under a quota.

Appendix B.2 notes that for $\alpha < \alpha^*$ the socially optimal emission price fluctuations are even stronger than the fluctuations under a quota. Moreover, in this case, a technological innovation reduces marginal abatement costs but increases socially optimal current emissions: the current innovation strongly reduces future abatement costs (and thus emissions) but only slightly reduces current abatement costs, making it optimal to emit more today in anticipation of the high reductions of future abatement costs.

4 The climate context

This section calibrates the dynamic model to reasonable climate change scenarios using DICE 2013 (Nordhaus 2013) and IPCC (2013). Despite our use of a stylized two-state model, the results illustrate the empirical relevance of the conceptual and theoretical insights described above: they may reverse the prices versus quantities ranking in plausible climate change scenarios.¹⁴

¹⁴Our model is consistent with an different interpretation that uses recent evidence of a near-linear relation between cumulative CO2 emissions and temperature change, the “transient climate response to cumulative CO2 emissions” (TCRE) (Matthews et al. 2009, MacDougall et al. 2017). With this alternative, we can set $\delta = 1$ and replace S , the carbon stock, with T , the temperature anomaly. The flow damage now depends on T instead of S . Recent climate economic applications of the TRCE model include Brock et al. (2014), Brock & Hansen (2017), and Dietz & Venmans (2018). This alternative requires a slightly different calibration than the one presented below.

4.1 Application

We use DICE 2013 to calibrate the abatement cost and damage parameters f and b . Our calibration uses annual values; in stating our units we suppress the “year”. Appendix A provides an explicit treatment of the time step. Setting abatement at 75% of the optimal level in DICE we obtain values for marginal abatement costs during the period 2015 - 2050. Fitting the average of these values to the linear marginal abatement cost function we obtain $f = 1.8$ in units of $\frac{G\$}{(GtCO_2)^2}$, our baseline value. We will contrast the resulting ranking with that obtained for a modified calibration that we refer to as *low abatement convexity*. We obtain this lower value of f by setting abatement at only 50% of the optimal level and averaging the marginal abatement costs of the DICE model over the period 2015-2100, resulting in $f = 1.2$ in $\frac{G\$}{(GtCO_2)^2}$.¹⁵

To calculate the damage parameter b we set annual Gross World Product (GWP) to the IMF’s 2016 estimate of 120 trillion in dollars using purchasing power parity.¹⁶ In DICE’s climate model, an increase of atmospheric carbon dioxide by 1270 $GtCO_2$ over the preindustrial level implies a medium to long-run temperature increase of $2^\circ C$. DICE assumes that this temperature increase lowers output by approximately 1%. This calibration assumption implies $b \approx 0.0015$ in units of $\frac{G\$}{(GtCO_2)^2}$, our baseline value.¹⁷

We also consider a *tipping or high damage convexity scenario* based on the concern that a temperature increase above $2^\circ C$ can trigger a variety of feedbacks that lead to a steep increase of damages. In this scenario we assume that the current temperature increase of approximately $1^\circ C$ causes negligible damage, but an increase of $3^\circ C$ causes a 5% loss in GWP. DICE’s climate model implies that an increase of 1600 $GtCO_2$ produces a $3^\circ C$ increase in temperature. If this increase lowers GWP by 5%, then $b \approx 0.0047$, our “tipping point” parameter estimate.¹⁸ We also consider a *low and flat marginal damage scenario* where we pick $b = 0.0005 \frac{G\$}{(GtCO_2)^2}$ for illustrative purposes.

We use the annual discount factor $\beta = 0.98$, consistent with the median

¹⁵The DICE 2013 marginal abatement costs change exogenously over time and endogenously in the emission level, hence the variation.

¹⁶This value is high relative to other estimates; the exact value in the World Economic Outlook Database, October 2017, is 120.197 trillion. We chose such a high value because our model (apart from h_t) is stationary, but world output is likely to grow substantially over the coming century. We note that DICE’s much lower estimate of world output is based on purchasing power parity weights that have recently been criticized for undervaluing output of the developing world (Deaton & Aten 2017)

¹⁷This conclusion uses $\frac{b}{2} (1271)^2 = 120 \times 10^3 (0.01)$. The factor 10^3 on the right side converts 120 \$Tr into 120×10^3 \$G(iga). The units of b are $\frac{G\$}{(GtCO_2)^2}$.

¹⁸This conclusion uses the calibration equation $\frac{b}{2} (1601)^2 = 120 \times 10^3 (0.05)$.

Table 1: Taxes versus Quantities Reassessed

meaning		base		tipping	low disc	low abate convexity	low/flat damages
technology diffusion	α	1	0.5	any	any*	any	(0.5,1)
damage convexity	b	0.0015		0.0047			0.0005
discount factor	β	0.98			0.99		
abatement convexity	f	1.8456				1.2	
preferred policy		tax	quota	quota	quota*	quota	tax

The table states the optimal policy instrument (final row) for different calibrations. The columns state the parameters that differ from the baseline calibration (first numeric column). Policy ranking holds for both annual and decadal time step except for the ‘low discount’ column (*); there, quotas dominate for full α domain only with a decadal time step; with an annual time taxes dominate if α is close enough to unity.

2% discount rate in the recent expert survey by Drupp et al. (2018, forthcoming). We also consider a high discount factor $\beta = 0.99$ as our *low discounting scenario*. We assume close to perfect persistence of an innovation, employing an annual correlation coefficient of 0.99. We calibrate the persistence of atmospheric carbon to Joos et al.’s (2013) model for carbon removal from the atmosphere, adopted in the 5th Assessment Report of the Intergovernmental Panel on Climate Change (IPCC 2013). A least square fit over 1000 years delivers an annual removal rate of 0.3% ($\delta = 0.997$), implying a half-life of 230 years.

Table 4.1 shows the policy ranking for our baseline scenario and its variations. The policy ranking is the same with either annual or decadal time steps except in the “low discount” column. There, quotas dominate for all α with a decadal time step, but taxes dominate if α is close to 1 and we use an annual time step.

For our baseline calibration with immediate technology diffusion ($\alpha = 1$) we conclude that taxes dominate quotas, in line with previous literature. However, reducing the diffusion parameter α from 1 to 0.5 causes quotas to dominate taxes. For the *tipping point* scenario, where we increase the damage convexity, quotas dominate for all values of α (including $\alpha = 1$, the feasible upper bound). Reducing the convexity of abatement cost or increasing

the discount factor both favor quotas. Our last column acknowledges recent literature that argues for relatively flat marginal damages, based on the closed-form integrated assessment model by Golosov et al. (2014). For very low and in particular flat marginal damages we get back to a scenario where taxes dominate quotas, at least for $\alpha \in (0.5, 1)$, i.e., sufficiently fast technology diffusion. From Proposition 2.ii we know that quotas dominate for sufficiently small α as long as the marginal damage curve is not entirely flat.

The previous sections explain why the widespread argument for taxes dominating quotas fails conceptually in the case of a stock pollutant. Our application demonstrates quantitatively that the case for taxes over quotas in the climate change context is much weaker than widely believed.

5 Conclusions

A widespread (static) criterion for ranking price-based and quantity-based regulation does not carry over to the dynamic setting where the regulated quantity is persistent. In this setting, the asymmetric information between the regulator and firms arises from technological change. The policy maker regulates an externality but does not observe recent innovations. The standard ranking criterion incorporates the effect of innovations on firms' cost structure. Our ranking criterion additionally incorporates the effect of the current innovation on firms' future production decisions and, thereby, externality cost arising from future damages. Both the persistence of the regulated quantity and the delayed technology diffusion favor quantity regulation.

Our discussion focuses on pollution control to mitigate climate change, where Weitzman's (1974) static ranking criterion is widely applied, even if only informally. However, contrary to the assumptions of Weitzman's model, all regulated greenhouse gases are persistent and the major greenhouse gas, carbon dioxide, persists for centuries. We emphasize that moving from flow to stock damages substantially increases damage convexity, i.e., the slope of the damage curve. We cannot judge the slope of the cumulative damage curve (the Social Cost of Carbon) based on the (generally very flat) annual damage curve.

Our main contribution is to derive a simple slope-based criterion for ranking prices versus quantities in the case of stock externalities under asymmetric information. Our graphical derivation furthers the intuition and produces an approximate ranking criterion. Our dynamic model formalizes the ranking criterion. There, we recognize that slope and shift parameters are endogenous. We also introduce a simple model of delayed technology diffusion and

demonstrate its policy relevance.

Our empirical application shows that the conceptual correction of the ranking argument substantially weakens the case for price regulation in climate change mitigation. We presented several reasonable calibrations for which cap and trade (quantity regulation) dominates taxes (price regulation). We selected our dynamic model to permit general analytic insight, restricting it to two state variables. As a result, the model remains a simple and stylized description of the complex assessment of climate change, even though we calibrate carefully to the integrated assessment literature and climate data. Our quantitative results do not imply that quotas necessarily dominate taxes in controlling carbon dioxide, but they demonstrate that our conceptual correction of the common ranking argument has serious policy implications.

Technological uncertainty lies at the heart of the Weitzman's (1974) asymmetric information problem. Technological change means that the regulator does not learn firms' current costs, even after many observations. In the pollution context, the long-lasting impact of current shocks on future abatement costs alters future emissions, changing social damages because these depend on cumulative emissions. Similar problems arise wherever asymmetric information is important and a regulator's objective depends on cumulative regulated actions.

References

- Aldy, J. E., Krupnick, A. J., Newell, R. G., Parry, I. W. & Pizer, W. A. (2010), 'Designing climate mitigation policy', *Journal of Economic Literature* **48**(4), 903–934.
- Allan, C., Jaffe, A. & Sin, I. (2014), 'Diffusion of green technology: a survey', Motu Working Paper 14-04.
- Brock, W., Engstrom, G. & Xepapadeas, A. (2014), 'Spatial climate-economic models in the design of optimal climate policies across locations', *European Economic Review* **69**, 78–103.
- Brock, W. & Hansen, L. P. (2017), 'Wrestling with uncertainty in climate economic models', Working Paper.
- Deaton, A. & Aten, B. (2017), 'Trying to understand the ppps in icp 2011: Why are the results so different?', *American Economic Journal: Macroeconomics* **9**(1), 243–264.

- Dietz, S. & Venmans, F. (2018), ‘Cumulative carbon emissions and economic policy: in search of general principles’, *Grantham Working Paper 283*.
- Drupp, M., Freemand, M., Groom, B. & Nesje, F. (2018, forthcoming), ‘Discounting disentangled: an expert survey on the determinants of the long term social discount rate’, *American Economic Journal: Policy*.
- Golosov, M., Hassler, J., Krusell, P. & Tsyvinski, A. (2014), ‘Optimal taxes on fossil fuel in general equilibrium’, *Econometrica* **82**(1), 41–88.
- Goulder, L. & Schein, A. (2013), ‘Carbon taxes versus cap and trade: A critical review’, *Climate Change Economics* **04**(03), 1350010.
- Hepburn, C. (2006), ‘Regulation by prices, quantities, or both: a review of instrument choice’, *Oxford Review of Economic Policy* **22**(2), 226–247.
- Hoel, M. & Karp, L. (2001), ‘Taxes and quotas for a stock pollutant with multiplicative uncertainty’, *Journal of Public Economics* **82**, 91–114.
- Hoel, M. & Karp, L. (2002), ‘Taxes versus quotas for a stock pollutant’, *Resource and Energy Economics* **24**, 367 – 384.
- IPCC (2013), *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Joos, F., Roth, R. & Weaver, A. J. (2013), ‘Carbon dioxide and climate impulse response functions for the computation of greenhouse gas metrics: a multi-model analysis’, *Atmos. Chem. Phys.* **13**, 2793–2825.
- Karp, L. & Zhang, J. (2005), ‘Regulation of stock externalities with correlated abatement costs’, *Environmental and Resource Economics* **32**, 273–299.
- Karp, L. & Zhang, J. (2012), ‘Taxes versus quantities for a stock pollutant with endogenous abatement costs and asymmetric information’, *Economic Theory* **49**, 371 – 409.
- MacDougall, A., Swart, N. & Knutti, R. (2017), ‘The uncertainty in the transient climate response to cumulative CO₂ emissions arising from the uncertainty in physical climate parameters’, *Journal of Climate* **30**, 813–27.
- Matthews, H., Gillett, N., Stott, P. & Zickeld, K. (2009), ‘The proportionality of global warming to cumulative carbon emissions.’, *Nature* **459**, 829–33.

- Newbery, D. (2018), ‘Policies for decarbonizing carbon in a liberalized power sector’, *Kiel Institute for the World Economy Economics Discussion Papers No 2018-29* .
- Newell, R. & Pizer, W. (2003), ‘Regulating stock externalities under uncertainty’, *Journal of Environmental Economics and Management* **45**, 416–32.
- Nordhaus, W. (2008), *A Question of Balance: Economic Modeling of Global Warming*, Yale University Press, New Haven. Online preprint: A Question of Balance: Weighing the Options on Global Warming Policies.
- Nordhaus, W. (2013), *The Climate Casino: Risk, Uncertainty, and Economics for a Warming World*, Yale University Press.
- Perino, G. & Requate, T. (2012), ‘Does more stringent environmental regulation induce or reduce technology adoption? when the rate of technology adoption is inverted u-shaped’, *Journal of Environmental Economics and Management* .
- Stavins, R. (1996), ‘Correlated uncertainty and policy instrument choice’, *Journal of Environmental Economics and Management* **30**(2), 218–232.
- Weitzman, M. (1974), ‘Prices versus quantities’, *Review of Economic Studies* **41**(4), 477–91.
- Weitzman, M. (2018), ‘Prices or quantities dominate banking and borrowing’, *NBER 24218* .
- Wood, P. J. & Jotzo, F. (2011), ‘Price floors for emissions trading’, *Energy Policy* **39**, 1746–1753.

A Appendix: Proofs

Preliminaries The intercept of the firms’ marginal benefit of emissions, $h_t + \rho\theta_{t-1} + \alpha\varepsilon_t$, incorporates exogenous changes via h_t . We treat the function h_t as deterministic, but nothing changes if we add to h_t an iid shock, uncorrelated with the technology shock, ε_t . If h_t is serially correlated, or if it depends on macro variables such as the business cycle, we would need to add additional state variables to the model. However, provided that those variables are uncorrelated with the technology shock and that they obey stochastic linear difference equations, this generalization does not alter the policy ranking.

We choose the unit of time to be one year and we use the parameter ϕ to represent the time step. Thus, if the time step is one decade, $\phi = 10$. This parameter does not appear in the formulae used in the text, because there we set $\phi = 1$. The parameter serves two purposes. First, it makes it simple to insure that the ratio used for policy ranking is unit-free. Second, the parameter enables us to calibrate the model to a particular time step, say one year, and then change the time step of the model without changing f or b .

For example, if the firms' benefit, during one year, of emitting at the annual rate of x_t is $(h_t + \rho\theta_{t-1} + \alpha\varepsilon_t)x_t - \frac{f}{2}x_t$, then their benefit of emitting at the same annual rate over a decade is $[(h_t + \rho\theta_{t-1} + \alpha\varepsilon_t)x_t - \frac{f}{2}x_t]10$. This formulation ignores discounting during a period. However, including intra-period discounting merely introduces a constant factor multiplying each period payoff, without changing the optimization problem or the policy ranking. Similarly, if the stock during a period is S_t , annual damages equal $\frac{b}{2}S_t^2$ and damages during a decade equal $\frac{b}{2}S_t^2 10$. This formulation again ignores intra-period discounting, and additionally assumes that the stock is constant during the period. In the climate context, the stock changes little during a year or a decade, so the assumption of a constant intra-period stock is unimportant. It would be easy to drop this assumption at the cost of slightly more complicated notation.

We take advantage of the linear-quadratic structure to avoid having to solve separate problems when the regulator uses taxes or quotas or in the first best (full information) setting. To this end, we introduce the indicator function

$$\Phi = \begin{cases} 1 & \text{if tax} \\ 0 & \text{if quota} \end{cases} .$$

We use $x_t \in \{e^T, e^Q\}$ to denote the regulator's control under tax and quantity regulation, respectively. With this notation, the regulator's problem, for $i \in \{T, Q\}$, is

$$\begin{aligned} \max \mathbf{E}_t \sum_{\tau=0}^{\infty} \beta \left[(h_{t+\tau} + \rho\theta_{t+\tau-1})x_{t+\tau} - \frac{f}{2}(x_{t+\tau}) + \Phi \frac{\alpha^2}{2f}\sigma^2 - \frac{b}{2}S_{t+\tau}^2 \right] \phi \\ \text{subject to } S_{t+\tau+1} = \delta S_{t+\tau} + \phi x_{t+\tau} + \Phi \phi \alpha \frac{\varepsilon_t}{f} \text{ and } \theta_t = \rho\theta_{t-1} + \varepsilon_t. \end{aligned}$$

The term $\Phi \frac{\alpha^2}{2f}\sigma^2$ in the payoff arises from taking expectations, in each period, of the shock for that period, ε_t . Here we use the assumption that these shocks are iid with mean zero. We refer to the problem formulated using x and Φ as the "generic problem" because it subsumes the problems under both taxes and quotas.

Because the problem has two state variables, it is convenient to use matrix notation. We define the state vector as $Y_t = (S_t, \theta_{t-1})'$ and we define the

following matrices:

$$Q = \begin{pmatrix} -b & 0 \\ 0 & 0 \end{pmatrix}, \quad A = \begin{pmatrix} \delta & 0 \\ 0 & \rho \end{pmatrix}, \quad W = (0 \quad \rho),$$

$$B = \begin{pmatrix} \phi \\ 0 \end{pmatrix}, \quad C = \begin{pmatrix} \Phi \phi \frac{\alpha}{f} \\ 1 \end{pmatrix}. \quad (8)$$

With this notation, the net flow payoff and equation of motion for the generic problem are:

$$\left[h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right] \phi \text{ and}$$

$$Y_{t+1} = A Y_t + B x_t + C \varepsilon_t.$$

Proof. (Lemma 1) The dynamic programming equation for the generic problem is:

$$J_t^i(Y_t) = \text{Max}_{x_t} \left[h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right] \phi + \beta \mathbf{E}_t J_{t+1}^i(Y_{t+1}). \quad (9)$$

The subscript t in J_t takes into account that the value function depends explicitly on calendar time due to the exogenous change in the intercept of marginal costs, h_t .

The value function for the LQ problem, for $i \in \{TQ\}$, is linear-quadratic: $J_t^i(Y_t) = V_{0,t}^i + V_{1t}' Y_t + \frac{1}{2} Y_t' V_2 Y_t$. The terms V_{1t} and V_2 are the same under taxes and quotas; only the term $V_{0,t}^i$ differs. The terms V_{1t} and $V_{0,t}^i$ inherit the time-dependence of h_t , but V_2 is constant. Denote $v_{1,t}$ as the first element of the column matrix V_{1t} , and define $\chi_t = -\beta v_{1,t}$, the intercept of the graph of the present value of the social cost of carbon. V_2 is a symmetric matrix, which we write as

$$V_2 = - \begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix}. \quad (10)$$

We write the difference in the payoff under taxes and under quotas as

$$\Delta_t \equiv V_{0,t}^T - V_{0,t}^Q.$$

Referees' Appendix B provides the details of the following steps:

1) We substitute the equations of motion into the right side of the DPE, equation 9, and take expectations.

2) We use the first order condition for x_t to obtain the linear control rule, $x_t = Z_{0t} + ZY_t$. The coefficients of the control rule, Z_{0t} and Z , are the same under taxes and quotas, a consequence of the ‘‘Principle of Certainty Equivalence’’; Z is a constant row vector and Z_{0t} is a time-varying scalar.

3) We substitute the optimal control rule back into the right side of the DPE to obtain the maximized DPE.

4) Equating coefficients of the terms that are quadratic in Y_t and independent of Y_t (on the two sides of the DPE) we obtain, respectively, an algebraic Riccatti equation for V_2 and a difference equation for V_{0t}^i .

This algorithm produces formulae for the endogenous parameters λ and μ . Using the definition $\varpi \equiv f \left(1 - \beta\delta^2 - \beta\frac{b}{f}\phi^2 \right)$, λ and μ satisfy

$$\lambda = \frac{1}{2\beta\phi} \left(-\varpi + \sqrt{\varpi^2 + 4\beta\phi^2bf} \right) > 0 \quad (11)$$

$$\mu = \frac{\lambda}{f} \left(\frac{\phi\beta\delta\rho}{(1 - \beta\delta\rho) + \beta\phi\frac{\lambda}{f}} \right). \quad (12)$$

From inspection of equation 11, $\lambda > 0$, so the numerator of the right side of equation 12 is positive. Therefore, μ has the same sign as ρ , which in our setting is positive, because the shock describes a technological innovation.¹⁹

We define $r \equiv \frac{b}{f}$, the ratio of the slopes of marginal damages and marginal benefit (equal to marginal abatement cost) and $R \equiv \frac{\lambda}{f}\phi$, the ratio of the slope of the SCC and the marginal flow benefit. The flexible time step ϕ enters the definition of R because we are interested in the ratio of the costs from an additional unit of emissions in the atmosphere λ and the benefits of emitting one more unit of emissions over the course of a period. If the period is not a year, then the benefit from *one* unit of emissions is $\frac{f}{\phi}$ rather than f . The parameter f measures the benefit from increasing the annual emission flow by one unit (so ϕ times the unit increase over the course of a period). Dividing both sides of equation 11 by f establishes Lemma 1. ■

Proof. (Proposition 1) Step 4 in the algorithm described in the proof of Lemma 1 also produces the difference equation for V_{0t}^i :

$$V_{0,t}^i = \left(h_t Z_{0t} - \frac{1}{2} f (Z_{0t})^2 + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi + \beta \left(V_{0,t+1}^i + V_{1t+1}' B Z_{0t} + \frac{1}{2} (B Z_{0t})' V_2 (B Z_{0t}) + \frac{1}{2} C' V_2 C \sigma^2 \right).$$

¹⁹The units of λ are $\frac{USD}{GtCO_2^2}$: The units of ϖ coincide with those of f and $\frac{1}{\phi}$ eliminates the time unit in f . The value function parameter μ is unit-free.

(Appendix B derives this relation; see equation 26.) We define $\Delta_t \equiv V_{0,t}^T - V_{0,t}^Q$, the difference in payoff under taxes and quotas. Using the fact that Z_{0t} , V_{t+1} , and V_2 are the same under taxes and quotas, and the definitions of Φ and C , we obtain the difference equation

$$\begin{aligned} \Delta_t &= V_{0,t}^T - V_{0,t}^Q = \frac{\alpha^2}{2f} \sigma^2 \phi + \beta \Delta_{t+1} \\ -\frac{1}{2} \beta \sigma^2 \left[\left(\begin{array}{cc} \phi \frac{\alpha}{f} & 1 \end{array} \right) \begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix} \begin{pmatrix} \phi \frac{\alpha}{f} \\ 1 \end{pmatrix} - \left(\begin{array}{cc} 0 & 1 \end{array} \right) \begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \right] \Rightarrow \\ \Delta_t &= \beta \Delta_{t+1} + \frac{\alpha^2}{2f} \sigma^2 \phi - \frac{1}{2} \beta \sigma^2 \phi \alpha \frac{\phi \alpha \lambda + 2\mu f}{f^2} = \beta \Delta_{t+1} + \frac{\alpha \phi}{2f} \sigma^2 \left(\alpha - \beta \frac{\phi \alpha \lambda + 2\mu f}{f} \right). \end{aligned}$$

The last line follows from carrying out the matrix multiplication and then simplifying. The steady state of this equation is the constant

$$\Delta = \frac{1}{1 - \beta} \frac{\alpha \phi}{2f} \sigma^2 \left(\alpha - \beta \frac{\phi \alpha \lambda + 2\mu f}{f} \right). \quad (13)$$

Using the definition $R \equiv \frac{\lambda}{f} \phi$ we have

$$\Delta = \frac{1}{1 - \beta} \frac{\alpha \phi}{2f} \sigma^2 (\alpha - \beta (2\mu + \alpha R)).$$

This equation implies that taxes dominate quotas if and only if

$$\alpha - \beta (2\mu + \alpha R) > 0. \quad (14)$$

Rearranging this inequality establishes inequality (7).

We rearrange inequality (14) using the definition of R and equation (12) for μ to obtain

$$\alpha - \beta \left(2 \frac{\beta \delta \rho R}{(1 - \beta \delta \rho) + \beta R} + \alpha R \right) > 0.$$

Multiplying by the positive denominator, this inequality is equivalent to

$$\begin{aligned} &\alpha ((1 - \beta \delta \rho) + \beta R) - \beta (2(\beta \delta \rho R) + \alpha R ((1 - \beta \delta \rho) + \beta R)) > 0 \\ \Leftrightarrow &-\beta^2 \alpha R^2 + (\alpha \beta - \beta (2\beta \delta \rho + \alpha (1 - \beta \delta \rho))) R + \alpha (1 - \beta \delta \rho) > 0 \\ \Leftrightarrow &R^2 - \frac{(\alpha \beta - \beta (2\beta \delta \rho + \alpha (1 - \beta \delta \rho)))}{\beta^2 \alpha} R - \frac{\alpha (1 - \beta \delta \rho)}{\beta^2 \alpha} < 0 \\ \Leftrightarrow &R^2 + \frac{1}{\alpha} \delta \rho (2 - \alpha) R - \frac{(1 - \beta \delta \rho)}{\beta^2} < 0 \\ \Leftrightarrow &R^2 + \kappa_1 R - \kappa_0 < 0, \end{aligned}$$

where the last inequality uses the definitions $\kappa_1 \equiv \frac{\delta\rho(2-\alpha)}{\alpha} > 0$ and $\kappa_0 \equiv \frac{1-\beta\delta\rho}{\beta^2} > 0$.

The quadratic expression $R^2 + \kappa_1 R - \kappa_0$ is negative at $R = 0$ and remains negative for R smaller than the positive root of the quadratic, defined as R^{crit} . Hence the inequality is satisfied for $R \in [0, R^{crit})$. ■

Proof. (Proposition 2) (i) In the first best (full information) world the regulator observes the technology shock in each period before choosing the level of emissions. Here, the regulator conditions emissions on S_t , θ_{t-1} and ε_t . Under asymmetric information and quotas, the regulator chooses emissions conditioned on S_t , θ_{t-1} and $E\varepsilon_t = 0$: under the quota, emissions do not depend on ε_t . Thus, the quota might be first best only if the first best level of emissions does not depend on ε_t .

We use well-known properties of the linear quadratic problem to show that the independence of the first best level of emissions and ε_t is sufficient, not merely necessary, for the quota to be first best. By the Principle of Certainty Equivalence for the linear quadratic problem, the coefficients of the linear and quadratic parts of the value function, V_{1t} and V_2 , are the same under taxes and quotas in the scenario with asymmetric information and also in the first best scenario. Thus, the parameters χ_t , λ , and μ are the same across the three scenarios.

The first best level of emissions equates the realized MAC and the present value of the social cost of carbon:

$$\rho\theta_{t-1} + \alpha\varepsilon_t - fE_t^{FB} = \beta (\chi_t + \lambda (\delta S_t + E_t^{FB}) + \mu (\rho\theta_{t-1} + \varepsilon_t)), \quad (15)$$

where E_t^{FB} denotes the first best level of emissions. An innovation ε_t causes the MAC curve to shift up by $\alpha\varepsilon_t$, and it causes the present value of the SCC to shift up by $\beta\mu\varepsilon_t$. We obtain the first order condition for the quota under asymmetric information by replacing ε_t with $E\varepsilon_t = 0$ and by replacing E_t^{FB} with E_t^Q (the quota) in equation 15. The fact that χ_t , λ , and μ are the same in the first best world and under quotas (and also under taxes) implies that the quota is first best if and only if the first best level of emissions does not depend on ε_t . From equation 15 this necessary and sufficient condition is equivalent to $\alpha = \beta\mu$.

Thus, to establish part (i) of the Proposition we need only establish that there exist an $\alpha \in (0, 1]$ that satisfies $\alpha = \beta\mu$. We have already established (for $\rho > 0$, our maintained assumption) that $\mu > 0$. To complete the proof we

need only confirm that $\beta\mu \leq 1$. Using the definitions of μ and R , we have

$$\begin{aligned} \beta\mu \leq 1 &\Leftrightarrow \beta^2 R \frac{\delta\rho}{(1-\beta\delta\rho)+\beta R} \leq 1 \Leftrightarrow \\ &R\beta(\beta\delta\rho - 1) \leq (1 - \beta\delta\rho). \end{aligned} \tag{16}$$

Because $\beta\delta\rho$ is bounded away from 1 and $R > 0$, the last inequality is always satisfied. Therefore, there exists $\alpha \in (0, 1]$ that satisfies $\alpha = \beta\mu$.

(ii) To show that a reduction in α favors quotas, we note that R^{crit} is a differentiable function of α . Using the chain rule and the definitions of κ_1 and κ_0 , we obtain

$$\frac{dR^{crit}}{d\alpha} = -\frac{1}{2} \frac{\kappa_1 - \sqrt{\kappa_1^2 + 4\kappa_0}}{\sqrt{\kappa_1^2 + 4\kappa_0}} 2\delta \frac{\rho}{\alpha^2} > 0. \tag{17}$$

Therefore, a reduction in α lowers the critical value R^{crit} , above which quotas dominate taxes.

To establish the second part of Part (ii), we note from Part (i) that for $\alpha = \beta\mu$ the quota is first best. Under the tax (using $E^T = e^T + \alpha \frac{\varepsilon_t}{f}$), we have

$$\frac{dE^T}{d\varepsilon_t} = \frac{\alpha}{f} > \frac{\alpha - \beta\mu}{f + \beta\lambda} = \frac{dE^{FB}}{d\varepsilon_t}, \tag{18}$$

where the second equality uses the first order condition 15 and the inequality uses $\lambda > 0$ and $\mu > 0$. This inequality means that emissions under the tax are always more responsive to a shock, compared to the first best level of emissions. Therefore, the tax can never support the first best level of emissions; quotas strictly dominate taxes for $\alpha = \beta\mu$, where the quota is first best. This fact and inequality 17 imply that quotas strictly dominate taxes for $\alpha \leq \alpha^* = \beta\mu$. The fact that this dominance is strict means that there exists $\alpha^{crit} > \alpha^*$ for which quotas strictly dominate taxes when $\alpha < \alpha^{crit}$. ■

B Referees' appendix

The first part of this appendix collects the details summarized by the algorithm in the proof of Lemma 1. The second part provides heuristic arguments for Propositions 1 and 2.

B.1 Material for Lemma 1

The equations of motion are

$$Y_{t+1} = AY_t + Bx_t + C\varepsilon_t$$

and the value function is

$$J_t(Y_t) = V_{0,t}^i + V_{1t}'Y_t + \frac{1}{2}Y_t'V_2Y_t.$$

The right side of the DPE, equation 9, is

$$\begin{aligned} & \left(h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi + \beta \mathbf{E}_t J_{t+1}(Y_{t+1}) \\ &= \left(h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi \\ & \quad + \beta \mathbf{E}_t \left(V_{0,t+1}^i + V_{1t+1}' Y_{t+1} + \frac{1}{2} Y_{t+1}' V_2 Y_{t+1} \right). \end{aligned}$$

Substituting in the equations of motion, we write the right side of the DPE as

$$\begin{aligned} & \left(h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi + \\ & \quad \beta \mathbf{E}_t [V_{0,t+1}^i + V_{1t+1}' (AY_t + Bx_t + C\varepsilon_t) \\ & \quad + \frac{1}{2} (AY_t + Bx_t + C\varepsilon_t)' V_2 (AY_t + Bx_t + C\varepsilon_t)]. \end{aligned}$$

Taking expectations gives

$$\begin{aligned} & \left(h_t x_t - \frac{1}{2} f x_t^2 + \frac{1}{2} Y_t' Q Y_t + W Y_t x_t + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi + \\ & \quad \beta [V_{0,t+1}^i + V_{1t+1}' (AY_t + Bx_t) + \\ & \quad \frac{1}{2} (AY_t + Bx_t)' V_2 (AY_t + Bx_t) + \frac{\sigma^2}{2} C' V_2 C]. \end{aligned} \tag{19}$$

The first order condition is

$$\begin{aligned} & (h_t - f x_t + W Y_t) \phi + \beta (V_{1t+1}' B + B' V_2 B x_t + B' V_2 A Y_t) = 0 \Rightarrow \\ & h_t \phi + \beta V_{1t+1}' B + (W \phi + \beta B' V_2 A) Y_t = (f \phi - \beta B' V_2 B) x_t = 0, \end{aligned}$$

which implies the control rule

$$x_t = \frac{1}{f \phi - \beta (B' V_2 B)} (h_t \phi + \beta V_{1t+1}' B + (W \phi + \beta B' V_2 A) Y_t)$$

or

$$x_t = Z_{0t} + ZY_t \text{ with} \quad (20)$$

$$Z_{0t} = \frac{h_t \phi + \beta V'_{1t+1} B}{f \phi - \beta (B' V_2 B)} \text{ and } Z = \frac{W \phi + \beta B' V_2 A}{f \phi - \beta (B' V_2 B)}.$$

Substituting the control rule into the expectation of the right side of the DPE (expression 19) gives the maximized right side of the DPE:

$$\begin{aligned} & [h_t (Z_{0t} + ZY_t) - \frac{1}{2} f (Z_{0t} + Y'_t Z') (Z_{0t} + ZY_t) \\ & + \frac{1}{2} Y'_t Q Y_t + W Y_t (Z_{0t} + ZY_t) + \Phi \frac{\alpha^2}{2f} \sigma^2] \phi + \\ & \beta (V_{0,t+1}^i + V'_{1t+1} (A Y_t + B (Z_{0t} + ZY_t))) \\ & + \beta \left(\frac{1}{2} (A Y_t + B (Z_{0t} + ZY_t))' V_2 (A Y_t + B (Z_{0t} + ZY_t)) + \frac{1}{2} C' V_2 C \sigma^2 \right) \end{aligned} \quad (21)$$

The terms that are quadratic in Y in expression 21 are

$$\frac{1}{2} Y'_t [(Q - f Z' Z + 2W' Z) \phi + \beta (A + BZ)' V_2 (A + BZ)] Y_t \quad (22)$$

Here we use the fact that $W Y_t = Y'_t W'$ (because both are scalars) so

$$W Y_t (Z Y_t) = Y'_t W' Z Y_t.$$

Now we use the fact that for any matrix H , $Y' H Y = Y' H' Y$ (because $Y' H Y$ is a scalar). Therefore $Y' H Y = \frac{1}{2} Y' (H + H') Y$. We can write any quadratic form as a symmetric quadratic. Using this fact, write

$$Y'_t W' Z Y_t = \frac{1}{2} Y'_t (W' Z + Z' W) Y_t.$$

Using this result we write the quadratic part of the right side of the maximized DPE, expression 22, as

$$\frac{1}{2} Y'_t [(Q - f Z' Z + W' Z + Z' W) \phi + \beta (A + BZ)' V_2 (A + BZ)] Y_t$$

Equating coefficients of the quadratic terms on the left and right sides of the maximized DPE gives

$$V_2 = [(Q - f Z' Z + W' Z + Z' W) \phi + \beta (A + BZ)' V_2 (A + BZ)] \quad (23)$$

We simplify the right side of equation 23 using the definitions of Z (equation 20) and V (equation 10) and the matrices defined in equation 8 to write

$$Z = \left(-\frac{\delta\beta\lambda}{f+\beta\lambda\phi}, \frac{\rho-\beta\mu\rho}{f+\beta\lambda\phi} \right).$$

With this result, performing the matrix manipulation on the right side of equation 23, and using equation 10, gives a recursive system of equations in λ , μ , and ν :

$$-\begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix} =$$

$$\begin{bmatrix} -\frac{1}{f+\beta\lambda\phi} (f\beta\lambda\delta^2 + b\beta\lambda\phi^2 + bf\phi) & -\beta\delta\frac{\rho}{f+\beta\lambda\phi} (f\mu + \lambda\phi) \\ -\beta\delta\frac{\rho}{f+\beta\lambda\phi} (f\mu + \lambda\phi) & -\frac{\rho^2}{f+\beta\lambda\phi} (-\phi\beta^2\mu^2 + \lambda\nu\phi\beta^2 + 2\phi\beta\mu + f\nu\beta - \phi) \end{bmatrix}$$

The equation for λ is

$$\lambda = \frac{1}{f+\beta\lambda\phi} (f\beta\lambda\delta^2 + b\beta\lambda\phi^2 + bf\phi) \Rightarrow$$

$$\beta\phi\lambda^2 + (f - f\beta\delta^2 - b\beta\phi^2)\lambda - bf\phi = 0$$

or

$$\beta\phi\lambda^2 + \varpi\lambda - bf\phi = 0.$$

The last line uses the definition $\varpi \equiv f \left(1 - \beta\delta^2 - \beta\frac{b}{f}\phi^2 \right)$. The positive root of this quadratic is

$$\lambda = \frac{1}{2\beta\phi} \left(-\varpi + \sqrt{\varpi^2 + 4\beta\phi^2bf} \right). \quad (24)$$

We know that the correct root is positive (so $-\lambda < 0$), because the negative root implies that the payoff grows arbitrarily large and positive as the stock of carbon becomes large. However, large carbon stocks result in high damages and a negative payoff.

The equation for μ is

$$\mu = \beta\delta\frac{\rho}{f+\beta\lambda\phi} (f\mu + \lambda\phi) \Rightarrow \mu = \beta\delta\rho\frac{\lambda}{f} \frac{\phi}{1 - \beta\delta\rho + \beta\frac{\lambda}{f}\phi}. \quad (25)$$

Collecting the terms in expression 21 that are independent of Y_t and equating these to $V_{0,t}^i$ produces the difference equation

$$V_{0,t}^i = \left(h_t Z_{0t} - \frac{1}{2} f (Z_{0t})^2 + \Phi \frac{\alpha^2}{2f} \sigma^2 \right) \phi + \quad (26)$$

$$\beta \left(V_{0,t+1}^i + V_{1t+1}' (BZ_{0t}) + \frac{1}{2} (BZ_{0t})' V_2 (BZ_{0t}) + \frac{1}{2} C' V_2 C \sigma^2 \right).$$

B.2 Heuristic arguments

Equations 15 and 18 rely on only the Principle of Certainty Equivalence for the linear quadratic problem. They do not require formulae for χ_t , λ and μ , and therefore provide the basis for a heuristic argument. As noted in the proof of Proposition 2, the fact that χ_t , λ and μ are the same in the three scenarios (full information and asymmetric information under a tax or a quota) imply that emissions in the three scenarios differ only if $\varepsilon_t \neq 0$, i.e. when the shock does not equal its expected value. For $\varepsilon_t = 0$ we have $E_t^{FB} = E_t^Q = E_t^T$. Emissions under the tax and in the first best scenario are linear in ε_t , with derivatives given in equation 18; of course the quota is independent of ε_t . Thus, for $\alpha \neq \alpha^*$,

$$\frac{dE^T}{d\varepsilon_t} = \frac{\alpha}{f} > \frac{\alpha - \beta\mu}{f + \beta\lambda} = \frac{dE^{FB}}{d\varepsilon_t} > \frac{dE^Q}{d\varepsilon_t} = 0. \quad (27)$$

The welfare cost of deviating from the first best level of emissions is quadratic in the deviation and symmetric around a zero deviation. The deviation between actual and first best emissions under the tax is

$$\left(\frac{\alpha}{f} - \frac{\alpha - \beta\mu}{f + \beta\lambda} \right) \varepsilon_t,$$

and the deviation between actual and first best emissions under the quota is

$$- \left(\frac{\alpha - \beta\mu}{f + \beta\lambda} \right) \varepsilon_t.$$

We need to consider two cases: $\alpha - \beta\mu > 0$ and $\alpha - \beta\mu < 0$; for $\alpha = \beta\mu$ we know from Proposition 2 that quotas dominate. For $\alpha - \beta\mu > 0$ the absolute value of the deviation under the quota exceeds the absolute value of the deviation under taxes for all $\varepsilon_t \neq 0$ (so taxes dominate quotas) if and only if

$$\left(\frac{\alpha - \beta\mu}{f + \beta\lambda} \right) > \left(\frac{\alpha}{f} - \frac{\alpha - \beta\mu}{f + \beta\lambda} \right).$$

Rearranging this inequality and using the definition of R produces the first equality in Proposition 1. For $\alpha - \beta\mu < 0$ the absolute value of the deviation under the quota exceeds the absolute value of the deviation under taxes (so taxes dominate quotas) if and only if

$$- \left(\frac{\alpha - \beta\mu}{f + \beta\lambda} \right) > \left(\frac{\alpha}{f} - \frac{\alpha - \beta\mu}{f + \beta\lambda} \right) \Rightarrow 0 > \frac{\alpha}{f}.$$

This inequality is never satisfied, so for $\alpha - \beta\mu < 0$ quotas dominate taxes, as established in the proof of Proposition 2.

Note that for $\alpha - \beta\mu < 0$ a larger current innovation (higher abatement costs) reduces the current first best level of emissions. When α is small, a positive current innovation means that abatement costs rise by a small amount, but expected abatement costs for future periods are expected to rise by a large amount. In this case, it is optimal to reduce current emissions (relative to the case where $\varepsilon_t = 0$) in anticipation of high future emissions.

B.3 Intertemporal correlation of adopted technology levels

This appendix derives the formula in Footnote 3.

We have $\theta_t = \rho\theta_{t-1} + \varepsilon_t$ and $\hat{\theta}_t = \rho\theta_{t-1} + \alpha\varepsilon_t$. Solving for θ_{t+j} and using the definition of $\hat{\theta}_{t+j}$ produces

$$\theta_{t+j} = \rho^{j+1}\theta_{t-1} + \sum_{s=0}^j \rho^s \varepsilon_{t+j-s} \quad (28)$$

$$\hat{\theta}_{t+j} = \rho \left(\rho^j \theta_{t-1} + \sum_{s=0}^{j-1} \rho^s \varepsilon_{t+j-1-s} \right) + \alpha \varepsilon_{t+j} \Rightarrow \quad (29)$$

$$\text{var}_t(\hat{\theta}_{t+j}) = \rho^2 \sum_{s=0}^{j-1} \rho^{2s} \sigma^2 + \alpha^2 \sigma^2 \text{ and } \text{cov}_t(\hat{\theta}_t, \hat{\theta}_{t+j}) = \alpha \rho^j \sigma^2 \quad (30)$$

Using the formula for covariance, we have

$$\text{corr}_t(\hat{\theta}_t, \hat{\theta}_{t+j}) = \frac{\alpha \rho^j \sigma^2}{\sqrt{\left(\rho^2 \sum_{s=0}^{j-1} \rho^{2s} \sigma^2 + \alpha^2 \sigma^2 \right) (\alpha^2 \sigma^2)}} \quad (31)$$

Using $\sum_{s=0}^{j-1} \rho^{2s} = \frac{\rho^{2j}-1}{\rho^2-1}$ to simplify the denominator of the previous expression produces

$$\text{corr}_t(\hat{\theta}_t, \hat{\theta}_{t+j}) = \frac{\alpha \rho^j \sigma^2}{\sqrt{\left(\rho^2 \frac{\rho^{2j}-1}{\rho^2-1} \sigma^2 + \alpha^2 \sigma^2 \right) \alpha^2 \sigma^2}} = \frac{\rho^j}{\sqrt{\left(\rho^2 \frac{\rho^{2j}-1}{\rho^2-1} + \alpha^2 \right)}}, \quad (32)$$

yielding the formula in Footnote 3.