Journal of Environmental Economics and Management 125 (2024) 102951

Contents lists available at ScienceDirect

Journal of Environmental Economics and Management

journal homepage: www.elsevier.com/locate/jeem

Taxes versus quantities reassessed[☆]

Larry Karp [a,*], Christian Traeger [b]

[a] Department of Agricultural and Resource Economics, UC Berkeley, United States of America
[b] Department of Economics, University of Oslo; ifo Institute for Economics Research, Norway

ARTICLE INFO

JEL classification:
Q00
Q50
H20
D80

Keywords:
Policy instruments
Pollution
Climate change
Taxes
Quantities
Regulation
Uncertainty
Cap and trade
Technology

ABSTRACT

The ongoing debate concerning the ranking of taxes versus cap and trade for climate policy begins with Weitzman's (1974) seminal slope-based criterion and concludes that taxes likely dominate quotas. We challenge this conclusion and the intuition behind it. Because technology shocks and pollution stocks are both persistent, a technology shock alters the intercepts of both the marginal damage and abatement cost curves. The ratio of these two intercept shifts is as important as the ratio of slopes in ranking policies. Technology innovations diffuse gradually, strengthening the importance of the ratio of intercept shifts. For plausible parameter combinations, quotas can dominate taxes.

1. Introduction

Following the Paris Climate Agreement, 88 countries considered implementing either a tax or a cap and trade system to regulate greenhouse gas emissions. By 2023, 47 national jurisdictions regulate over 23% of global greenhouse gas emissions by an emission tax or an emission trading scheme (World Bank, 2023). At 2020 prices, the European Emissions Trading System alone has an approximate annual market value of over 150 billion USD. Both taxes and cap (quotas) and trade are second-best policy instruments when we face uncertain technological progress in abatement costs and macroeconomic shocks. Our paper provides new intuition for ranking taxes and cap and trade and challenges the widely accepted view that taxes dominate quotas in the climate context.

Weitzman's (1974) model of flow pollutants provides the basis for current intuition about the welfare ranking of taxes and quotas. Under a tax, uncertainty about marginal abatement costs creates uncertainty about emissions, and therefore about damages. Under a binding quota, the randomness creates cost uncertainty. Weitzman (1974) shows that taxes create a smaller deadweight loss than quotas if and only if the slope of marginal damages is less than the slope of marginal abatement costs. His paper is among the most widely taught in environmental economics, with over 4000 citations.

☆ We are grateful for comments from two reviewers and the Journal editor Andreas Lange and earlier comments by Lint Barrage, Inge van Den Bijgaart, Reyer Gerlagh, Christian Gollier, Rolf Golombek, Bard Harstad, Michael Hoel, Fortunat Joos, Eric Naevdal, Armon Rezai, Hiroaki Sakamoto, Robert Stavins, Jean Tirole, Martin Weitzman, and seminar participants at the EAERE 2017, SITE 2017, CESifo area network conference 2017, 2018 Workshops on Climate Change in Bologna and Copenhagen, SURED 2018, and EEA 2018. Christian Traeger gratefully acknowledges financial support by the Norwegian Research Council grant No. 315878.

* Corresponding author.
E-mail addresses: karp@berkeley.edu (L. Karp), christian.traeger@econ.uio.no (C. Traeger).

https://doi.org/10.1016/j.jeem.2024.102951
Received 22 March 2023
Available online 29 February 2024

Most pollutants, and all climate-related pollutants, have some persistence or cumulative impact; carbon dioxide's effective half-life exceeds a century. The literature has long recognized that for stock pollutants, the relevant marginal damage curve involves the discounted stream of future marginal damages arising from the changes in future pollution stocks caused by current emissions. When no confusion arises, in the stock pollutant context we refer to this discounted stream simply as "marginal damages"; when there is a possibility of ambiguity, we refer to it as the social cost of carbon (SCC).

The literature has largely transferred Weitzman's logic, developed for a flow pollutant, to the context of stock pollutants, merely substituting the slope of flow marginal damages with the slope of the discounted stream of marginal damages. It is generally agreed that this slope is very small. This widespread view on relative *slopes* (which we do not challenge), together with Weitzman's logic, has led to the conclusion that taxes dominate quotas for controlling climate change.[1] We explain why Weitzman's logic fails for climate change, and we derive a ranking criterion applicable to greenhouse gases. A calibration for carbon dioxide illustrates the quantitative importance of the difference.

We focus on the cost uncertainty arising from shocks to technology (deviations from a trend). Section 3 explains this modeling choice. The mechanism is symmetric with respect to positive and negative shocks, so we use the example of a shock that unexpectedly lowers marginal abatement costs in the current period. The optimal level of pollution responds to a change in abatement costs. Equilibrium abatement does not respond to this change under a binding quota, and the response is excessive under a tax. In both cases, the deviation between the equilibrium and the optimal response creates a deadweight loss. Weitzman shows that for a flow pollutant the deadweight loss is larger under the quota if the environmental damages are less sensitive than the costs to a change in emissions: taxes welfare-dominate quotas if the marginal damage curve is flatter than the marginal abatement costs curve. For climate change, there is wide agreement that the slope of marginal damages is smaller than the slope of marginal costs.

However, because technology is persistent, the technological innovation also reduces future costs. The cost shocks are firms' private information in the current period, but the regulator learns them in the next period by observing the equilibrium level of emissions in response to a tax, or the equilibrium permit price in the case of a quota. Either piece of information enables the regulator to infer the shock. The shock-induced reduction in future abatement costs reduces future emissions,[2] thereby lowering the stock trajectory and lowering the stream of future marginal damages: the SCC. As a result, the current shock shifts the SCC in the same direction as the marginal abatement cost curve. Thus, the persistence of *both* the technology and the pollution stock cause the marginal cost and the SCC to be perfectly positively correlated. This perfect correlation arises endogenously from the very nature of the stock pollutant problem.

To the best of our knowledge, our companion paper, Karp and Traeger (2024), is the first to recognize this positive correlation; the mechanism that creates it is absent if either the pollutant or the cost shock are not persistent. Karp and Traeger (2024) explores the implications for a new first best policy instrument, a "smart cap". The current paper explains how this mechanism alters the logic behind the policy-ranking, thereby favoring quotas.

We derive a new, simple ranking criterion for stock pollutants. Like Weitzman's criterion for flow pollutants, the policy ranking depends on the relative slopes of the marginal abatement cost curve and the marginal damage curve, provided we understand that the marginal damages involve the discounted cost stream. We show that these damages are much more convex than the corresponding flow damages; however, this increase in convexity alone does not cause quotas to dominate taxes. The additional effect is that the ranking for stock pollutants depends on the relative shifts in intercepts. This "intercept effect" is of first order for ranking taxes versus quotas. For plausible parametrizations, "intercept shifts are as important as slopes". Therefore, the accepted view that the slope of the SCC is very small does not imply that taxes dominate quotas.

Our focus on technology shocks brings into play a previously missing consideration: technology diffuses gradually. Gradual diffusion of technology is distinct from the persistence of technology. Gradual diffusion means that (for example) a cost-reducing shock may shift the current industry's marginal abatement costs down by only a small amount because only a small fraction of the industry adopts the technology. However, future marginal costs fall more substantially with high future adoption. The larger fall in future costs reduces future emissions further, thereby leading to a more substantial fall in the SCC.

That is, gradual diffusion makes the shock-induced shift in SCC larger relative to the shift in the current marginal abatement costs. This effect favors quotas because it further reduces the optimal adjustment of abatement. We use the stylized model of technology diffusion introduced in Karp and Traeger (2024). There, the speed of technology diffusion affects the characteristics of two first best policies, which we refer to as the smart cap and the smart tax. In the current paper, the speed of diffusion affects the ranking of the two most important second-best policies, the tax and cap and trade.

Our calibration uses estimates of marginal abatement costs and marginal damages based on Nordhaus and Sztorc (2013) and models climate change based on the "transient climate response to cumulative carbon emissions" (TCRE). The TCRE model posits a linear relation between cumulative emissions and temperature change and is extensively discussed in the IPCC (2013). Recent applications in economics include Anderson et al. (2014), Brock and Xepapadeas (2017), Dietz and Venmans (2018), and Dietz et al. (2021). The model avoids the exaggerated lag in warming generated by the DICE model (Dietz et al., 2021). Our calibration mostly follows Karp and Traeger (2024), adding additional scenarios. The full model is transparent, and produces a policy ranking

[1] Nordhaus (2008) writes "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)". Wood and Jotzo (2011) write "It is generally thought that [Weitzman's logic holds] ...with climate change for the comparison between price and quantity instruments". Weitzman (2018) writes "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".

[2] The same reasoning holds even under business as usual emissions without a regulator; a clean innovation will reduce future emissions.

that depends on only a few parameters. There is substantial agreement about (or at least familiarity with) all of these parameters except for the rate of technological diffusion, for which we conduct sensitivity studies. For some parameter sets we reproduce the conventional view that taxes dominate quotas, but for other plausible parameters, the ranking favors quotas.

We consider the problem of a global planner who, in each period, chooses a global quota or a common tax for all regions. Provided that there are technological spillovers, the qualitative problem is similar even if only a subset of countries institutes policies. Those countries will behave as our model describes, and the other countries will free-ride; however, the current technological innovation will affect future emissions in both sets of countries. In contrast, it might be rational for a noncooperative entity that operates on a smaller scale (e.g., one who chooses a policy that applies to a single sector) to ignore the dynamic or spatial externalites.

A large literature extends Weitzman's analysis for flow pollutants.[3] Most of this literature assumes that shocks affecting marginal costs and marginal damages are uncorrelated, making the marginal damage shocks irrelevant to the comparison between taxes and quotas. However, Weitzman (1974) recognizes that the correlation between separate shocks to marginal costs and damages complicates his original ranking criterion. Stavins (1996) elaborates on this point, providing different examples of exogenous correlations in the flow pollution setting, and developing the ranking criterion for arbitrary correlation. Thus, the fact that correlation between shifts in marginal abatement costs and damages has an important effect on the welfare ranking is well understood.

We make a different point, showing that a perfectly positive correlation between the marginal abatement costs and the SCC results endogenously from the very nature of the greenhouse gas abatement problem, for the reasons explained above. Because this correlation is perfect, we can develop a simple ranking criterion based on relative slopes and relative intercept shifts. These two ratios are endogenous for a stock pollutant; determining the ratios and then showing how they affect the welfare ranking requires the solution of a dynamic problem. In contrast, with a flow pollutant the correlation between marginal abatement costs and damages is given exogenously.

Our use of the TCRE model for temperature change, together with our adoption of the nearly uncontested view that damages are convex in temperature, implies that damages are convex in the stock of CO_2; therefore, the slope of the SCC is positive. That convexity assumption is widespread but not universal.[4] Our qualitative results hold even if damages are only slightly convex in the stock, as in our leading calibration.

Montero (2002) considers policy ranking under incomplete enforcement, and Shinkuma and Sugeta (2016) considers the ranking with endogenous firm entry. Keohane (2009) notes that the distribution of quota rents may make it easier to achieve political buy-in from regulated industries under quotas rather than taxes. Requate and Unold (2003) show how the incentives to adopt technology vary with the instrument choice, and Perino and Requate (2012) show how policy stringency alters these incentives. Pizer and Prest (2020) use a two-period model with linear damages. Recent reviews of taxes and quotas that discuss both stock and flow pollutants include Hepburn (2006), Aldy et al. (2010), Goulder and Schein (2013), Newbery (2018), and Stavins (2020).[5]

A smaller literature compares policies for a stock pollutant. Pizer (1999) and Fischer and Springborn (2011), and Heutel (2012) use simulations to compare policies under uncertainty, but where firms and the regulator have the same information. Heutel (2012) also briefly considers the role of asymmetric information between firms and the regulator. The latter two papers focus on the effect of business cycles on optimal policy.

A separate literature extends Weitzman's linear-quadratic asymmetric information model to produce analytic comparisons for a stock pollutant. In this setting, there is an important difference between open loop and feedback policies. In the former, a regulator at t chooses the sequence of current and future policy levels conditional on information available at t. With feedback policies, a regulator at t chooses the current policy level and understands that future policy levels will be conditioned on information that becomes available in the future. Hoel and Karp (2002) assume serially uncorrelated shocks, thereby ruling out our results. Newell and Pizer (2003) consider serially correlated cost shocks, but only in an open loop setting. They show that positively correlated cost shocks increase stock volatility under taxes, favoring quotas.[6] Karp and Zhang (2005) compare the policy ranking across the open loop and feedback settings, and find that positive serial correlation of cost shocks favors quotas under feedback policies.[7]However, they do not explain the mechanism or include our model of gradual diffusion, and they do not discuss technology.

This literature shows that increasing the regulator's ability to respond to information, either by moving from the open loop to the feedback setting, or by reducing the time step between policy adjustments within the feedback setting, *both favor taxes*. For a long-lived problem such as climate change, a policymaker understands that future policies adjust to new information. Therefore,

[3] Most of this literature is partial equilibrium. The exception, Kelly (2005), uses a general equilibrium model, where consumers' risk aversion favors quotas.

[4] van der Ploeg et al. (2022) assume linearity of damages in temperature based on an estimate by Kalkuhl and Wenz (2020). A somewhat different strand of literature builds on Golosov et al.'s (2014) argument that the convexity of damages in temperature can be offset by the concavity of temperature in the CO_2 concentration, leading to a potentially linear welfare impact of CO_2 emissions (but not linearity of damages in temperature). Traeger (2023) develops this model into a full-fledged IAM with temperatures closely matching the CMIP5 data.

[5] Montero (2008) and Boleslavsky and Kelly (2014) show how the regulator can use mechanisms or the timing of policy changes to induce firms to truthfully reveal their abatement costs, thereby eliminating the asymmetry of information. Most real-world policies are not designed to elicit information. We compare the two simplest policies, taxes and cap and trade, where firms do not behave strategically.

[6] Under open loop taxes, positively serially correlated shocks produce 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 open loop quotas, 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.

[7] Karp and Zhang (2006) show that anticipated learning about climate-related damages favors taxes, but the effect is small. Karp and Zhang (2012) study policy ranking with endogenous investment in abatement capital. Neither paper includes persistent technology shocks.

we consider only the feedback setting. However, we leave the time step as a model parameter, and our application assumes that policy adjusts every five years. Our results do not rely on rapid adjustment of policies.

None of the papers that formally rank taxes and quotas for stock pollutants include our conceptual insight explaining why Weitzman's reasoning does not carry over to stock pollutants with persistent shocks. This insight can be conveyed in a figure nearly as simple as Weitzman's graph, and therefore can be taught at the undergraduate level (Fig. 1). The earlier literature also misses our other contributions: We provide a simple and intuitive ranking criterion which shows that the optimal policy choice depends as much on the ratios of intercept shifts as on ratios of slopes of marginal abatement costs and damages. We provide empirical evidence that the case for using taxes instead of cap and trade as a climate policy is weaker than previously thought. Finally, we show that quotas can even be first best, despite an almost flat marginal damage curve, if the intercept effect is strong enough.

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 fundamental 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, E: $MD = a + gE$. The slope parameter g 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.[8] The upper left panel in Fig. 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.[9] With taxes, the policy fixes the emissions price at the green (horizontal) line in Fig. 1. In a quantity setting, the policy caps the emissions at the red (vertical) line.

The dashed lines in Fig. 1 shows the realized marginal abatement cost curve for a shock that reduces marginal abatement costs. The results are symmetric with respect to positive and negative shocks, so we illustrate only $\theta < \mathbb{E}(\theta)$. The top left panel shows the tax and the quota equilibria for a flow pollutant. The deadweight loss under the tax is the light green triangle and the deadweight loss under a quota is the heavy red triangle. The deadweight loss is smaller under the tax than under the quota because the MAC curve is steeper than the MD curve in this figure. Here, taxes dominate quotas.

2.2. Modification for a stock pollutant

With a flow pollutant, the relevant MD is the current marginal damage arising from an additional unit of pollution today. The top panels of Fig. 1 incorporate Weitzman's (1974) assumption that θ does not shift these *flow* marginal damages. We maintain the same assumption in the dynamic model with stock pollutants. The relevant marginal damage is the present discounted *stream of future marginal damages*, commonly known as the Social Cost of Carbon (SCC): the change in the expectation of all future climate-related damages arising from an increase in the current flow of pollution.

The flow marginal damage is primitive, but the SCC is an endogenous function because the additional discounted stream of costs arising from an extra unit of emissions depends on future emissions. Section 3 derives the SCC; here, in the interest of simplicity, we take it as given. The lower panels in Fig. 1 show the graph of this function, which we again label MD (in order to emphasize the link between the top and bottom panels). The reader should keep in mind that MD in the lower panels refers to the SCC.

The similarity between the top and the bottom panels has created the impression that we can rank policies for a stock pollutant using the same criterion as for a flow pollutant, merely replacing the slope of the flow MD with the slope of the SCC. However, as explained below, a persistent technology shock shifts the MAC and the SCC in the same direction. The bottom panels in Fig. 1 illustrate this shift. It is evident from these panels that the perfect correlation between the MAC and the SCC erodes (and possibly even reverses) the advantage of taxes — despite the fact that the slope of the SCC is smaller than the slope of the MAC.

---

[8] We emphasize that the marginal benefits from emissions are equal to the marginal abatement costs. Let 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.

[9] The common assumption that the intercept but not the slope is private information is key to the simplicity of both Weitzman's and our result. Hoel and 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.

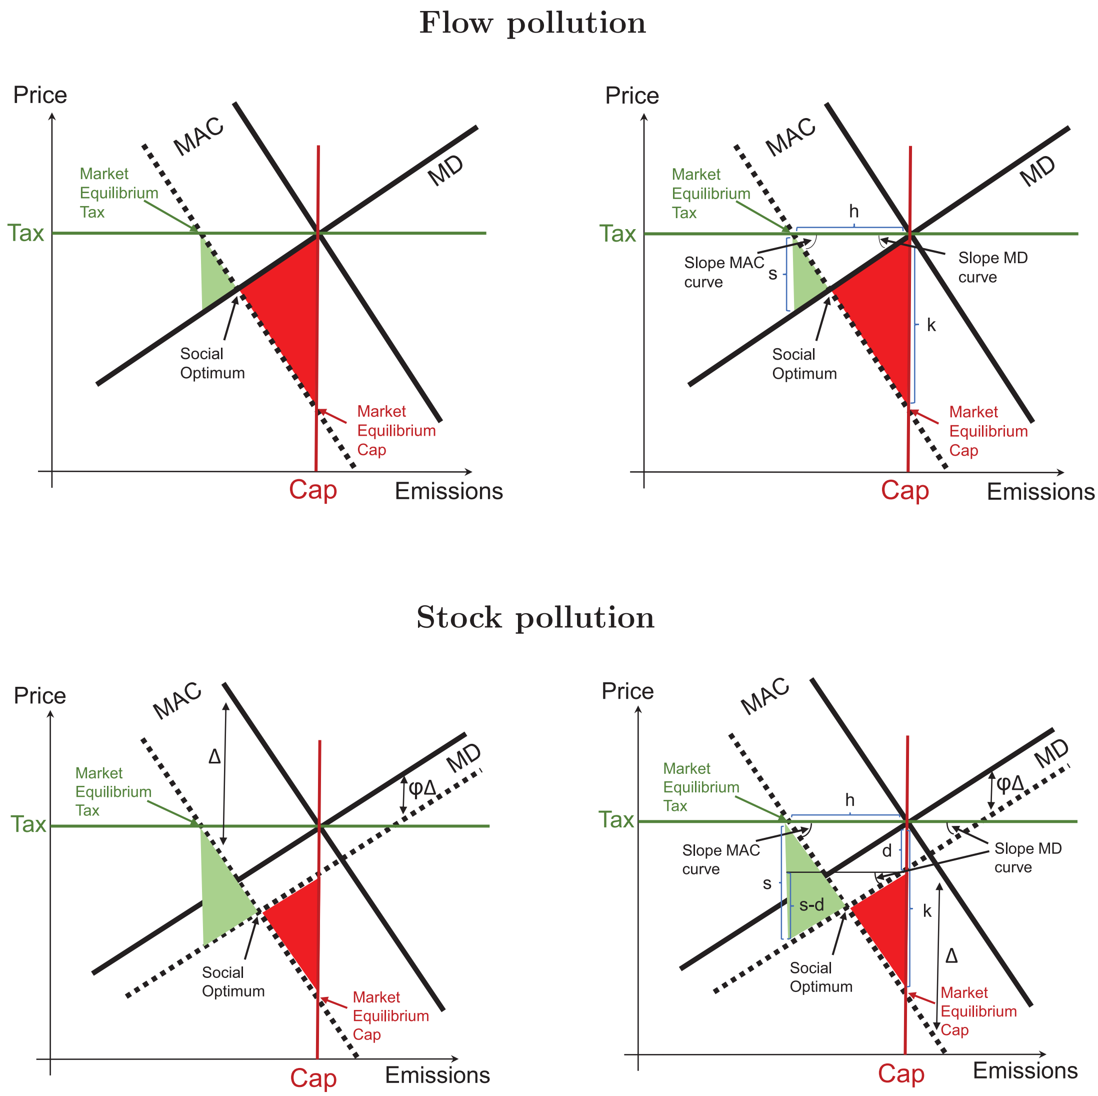

Fig. 1. Illustration of Weitzman (1974) insights for a flow pollutant (top panels) and a quasi-static illustration of the changes for a stock pollutant (bottom panels). The light green (left) triangle characterizes the deadweight loss under a tax, whereas the red (right) triangle characterizes the deadweight 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..

We now explain this correlated shift, and then return to the figure to quantify its effect. Technology is persistent. An unexpectedly successful innovation not only lowers today's abatement costs, but also future abatement costs. As a result, future equilibrium emissions fall.[10] Thus, the shock that lowers today's marginal abatement costs also lowers future CO_2 concentrations. Lower CO_2 concentrations in the future imply lower marginal damage from today's emissions and a reduction of the SCC. Analogously, lower than expected technological progress increases the marginal abatement costs today and in the future, and it increases future CO_2 concentrations. The higher CO_2 concentrations imply higher than expected marginal damages from today's emissions and, thus, a higher SCC.

The lower left panel in Fig. 1 illustrates the consequences of this insight. The two dashed curves show MAC and MD conditional on Δ, the shift in marginal abatement cost. The parameter φ is the ratio of the shift in the intercept of MD to the shift in the MAC curve. Our graphical analysis takes the slope of the MD curve and the ratio of shifts in intercepts (φ) as exogenous, and also assumes that these are the same under taxes and quotas, for any realization of the cost shock, Δ. Our genuinely dynamic model (Section 3)

[10] We assume throughout that future policy is set optimally. However, even in the absence of any regulation, a persistent shift in the marginal abatement cost changes the trajectory of future BAU emissions, causing a shift in the SCC. The magnitude but not the direction of the shift depends on our assumption that future policy is set optimally.

recognizes that both the slope of the MD curve and ratio of intercept shifts are endogenous functions of the model parameters. However, these two functions are the same under taxes and quotas, and they are independent of the cost shock. This invariance is important, because without it the static model in this section would shed little light on the dynamic model that we use to study a stock pollutant. Section 2.3 explain this invariance.

For our example in the Figure, an innovation lowers, by Δ, the marginal abatement cost from the solid line to the dashed line. Because technology is persistent, this reduction in marginal abatement costs makes future emission reductions cheaper, and reduces future emissions (both optimal and Business-as-Usual, BAU).[11] The resulting reduction in the future trajectory of pollution stocks lowers the marginal damage from releasing an additional unit of the pollutant today. As a consequence, the MD curve also shifts by $\varphi\Delta$.

Comparing the top and the bottom left panels in Fig. 1, it is apparent that the downward shift of the MD curve increases the deadweight loss of the tax and reduces the deadweight loss of the quota, thus favoring quotas. The right panels in Fig. 1 enrich the left panels by adding labels to some slopes and segment lengths. Using these labels, a familiar geometric argument establishes Weitzman's result that taxes dominate quotas for a flow pollutant if and only if the slope of the marginal damage curve is less than the slope of marginal abatement cost curve. A similar argument (Appendix B) shows that taxes dominate quotas for a stock pollutant if and only if

$$\frac{g}{f} \equiv \frac{m^{MD}}{m^{MAC}} < 1 - 2\varphi \tag{1}$$

The ranking of taxes versus quantities now depends on both the ratio of the slopes of the two curves and on the ratio of their shifts, φ. The figure shows that a shock induces a greater than optimal emissions adjustment under a tax, and a less than optimal adjustment under a quota. The deadweight loss is monotonic in the deviation between the equilibrium adjustment and the socially optimal adjustment: taxes dominate quotas if and only if the deviation is greater under quotas than under taxes. A shift of the MD curve does not alter the equilibrium emissions adjustment under taxes or quotas, but it reduces the socially optimal adjustment, moving it closer to the equilibrium under quotas (no adjustment). Therefore, a positive value of φ always lowers the deadweight loss under quotas and raises the deadweight loss under taxes.

2.3. Discussion of the graphical model

A footnote in Weitzman (1974), elaborated by Stavins (1996), considers the case where θ and a shock that shifts the MD curve are correlated. Stavins gives the example where 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. Here, the correlation between the shocks affecting marginal abatement costs and damages is a primitive, taking any value in $[-1, 1]$.

In contrast, the correlation between marginal abatement costs and damages (that is, the SCC) with a stock pollutant is endogenous, arising from the future response of emissions to a current cost shock. Moreover, correlation is perfect (correlation coefficient of unity). Identifying the correlation here is trivial, but the key to the ranking is φ, the ratio of the shift in the marginal damage curve per unit shift in the marginal abatement cost curve. Given that the mechanisms for the correlation under flow and stock pollutants are entirely different, it is perhaps not surprising that it took a quarter of a century after Stavins' paper to make the link between the two problems.[12]

Our graphical treatment takes the two components of the ranking criterion, the ratio of slopes and the ratio of intercept shifts (φ) as exogenous. These are, of course, endogenous objects in the dynamic setting. Our graphical treatment also assumes that these two ratios do not depend on whether the regulator uses a tax or quota. This invariance is a result, not an assumption, in the dynamic model in Section 3. This fact is important because it means that our graphical treatment accurately reflects the forces at work in the genuinely dynamic model. The dynamic model makes it possible to calculate the values of the two ratios and thereby rank the tax and quota for a stock pollutant.

The invariance is due to three implications of the linear-quadratic model with additive shocks. (i) The value function is quadratic in the pollution stock and the technology shock. This fact means that the realization of the SCC is a linear function of the next-period pollution stock and the technology shock. (ii) The coefficients of the linear and quadratic terms in the value functions under optimal taxes and quotas are identical.[13] This fact means that both the intercept and the slopes of the SCC are the same under the two policies. Therefore, conditional on the next-period stock, the SCC is the same under taxes and quotas. (iii) The *expected* level of emissions, and therefore the expected next-period pollution stock, is the same under the optimal tax and quota. This fact, together with the linearity (in pollution stock) and the equivalence (under taxes and quotas) of coefficients, means that the SCC in the genuinely dynamic model is invariant to the choice of policy — exactly as our static model assumes.

[11] We assume throughout that future policy is set optimally. However, even in the absence of any regulation, a persistent shift in the marginal abatement cost changes the trajectory of future BAU emissions, causing a shift in the SCC. The magnitude but not the direction of the shift depends on our assumption that future policy is set optimally.

[12] The effect of the positive correlation (unlike the mechanism behind it) is the same in the two settings: positive correlation favors quotas.

[13] In contrast, the intercepts of the two value functions differ. The policy ranking depends on a comparison of those intercepts.

3. The dynamic model

Two sources of asymmetric information cause the non-equivalence of taxes and quotas in the dynamic setting. First, asymmetry arises because technology-related costs are private information when firms choose emissions. Second, asymmetry arises because emissions decisions occur more frequently than policy updates, unless the regulator can condition the policy instrument on the arriving public information.

Our analysis focuses on the asymmetric information resulting from technological innovation, which we consider most relevant for three reasons. First, many technological innovations are genuinely private or unverifiable information at the time firms choose emissions. Second, it is hard to condition policy on technological innovation. Third, technological innovations are persistent, and therefore affect all future periods. The impact of technology shocks is therefore not easily mitigated by intertemporal arbitrage. Many macro-economic shocks dissipate over five years (the length of a period in our calibration) although some are more persistent. Our model easily accommodates such persistence (Online Appendix C.1). Intertemporal arbitrage through banking and borrowing of certificates, or through the European Emission Trading System's market stability reserve, reduce market disturbances from such shocks. Moreover, conditioning carbon policy on macroeconomic indicators can eliminate much of the costs arising from macroeconomic shocks (Ellerman and Wing, 2003; Jotzo and Pezzey, 2007; Newell and Pizer, 2008; Doda, 2016; Burtraw et al., 2020). (Appendix C.1).

Policy makers will acquire new information during the many decades that climate policy remains relevant. They will adapt regulation following unexpected changes in the cost of renewable energy generation. Even if current policy makers do not intend to adapt policy in the future, they cannot commit their successors to ignore new information for long periods of time. Pizer and Prest (2020) note that "most real-world regulations are updated over time in response to new information". We therefore study a model where the policy maker in each period conditions current regulation on current information, and understands that future policy makers will do the same: the policy rule is feedback, not "open loop with revision". Online Appendix C.4 shows that the feedback policy can be implemented by announcing state-contingent policy rules at the beginning of the planning horizon.

3.1. Description of the model

The model contains two state variables: a pollution stock and a technology level. The equation of motion for the pollution stock is

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

where E_t are emissions. The classic stock pollution model interprets S_t as the pollutant's concentration and $1-\delta \geq 0$ as the pollutant's decay rate. For our climate change application, we use the fact that atmospheric temperature increase is approximately proportional to cumulative historic emissions and we interpret S_t as temperature (and therefore set $\delta = 1$; see Section 4 for details). The stock S_t causes annual damages of $\frac{b}{2}S_t^2$.[14] The exogenous parameter b equals the slope of the marginal *flow* damage curve. Both the (discounted stream of) the marginal damage from releasing another unit of emissions and its dependence on technology shocks are endogenous to the model, not an exogenous input as in Section 2.

The abatement technology consists of a deterministic trend and a stochastic deviation θ_t from this trend. This deviation is a highly persistent stochastic process under iid shocks $\varepsilon_t \sim iid\left(0, \sigma^2\right)$. These shocks represent technological innovations departing from the trend. The equation of motion for θ (a component of the technology stock) 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 choosing the policy for period t; firms know both θ_{t-1} and ε_t in period t. This asymmetry provides the dynamic analogue of Weitzman's (1974) asymmetric information.

The speed of technological diffusion plays a critical – and a novel – role in ranking the policies. A large literature documents the fact that many new technologies diffuse gradually through the economy (Rogers, 2003). The standard approach uses an S−shaped function to describe the relation between the time since a new technology was introduced and the fraction of firms that have adopted it. We use a simpler model, introduced in Karp and Traeger (2024), in which the representative firm adopts only the fraction $\alpha \in (0, 1]$ of the latest technological innovation during the current period, and adopts the remaining fraction in the next period. This model captures gradual (exogenous) diffusion without the need of an additional state variable.[15] We define $\hat{\theta}_t \equiv \rho\theta_{t-1} + \alpha\varepsilon_t$ as the stochastic component of the *adopted* technology. We let h_t denote the deterministic trend of adopted technology; this trend is important for our calibration but it does not appear in the formula for policy ranking. Our formulation complements an AR(1) model for *innovated technology* (θ) by an ARMA(1,1) structure for *adopted technology* ($\hat{\theta}$).[16]

The persistence of innovated technology implies a high autoregressive coefficient ρ, and thus high serial correlation for adopted technology. A lower value of α further increases this serial correlation because a given level of technology adoption today results

[14] The absence of a linear damage term results from defining S_t as the deviation of the stock from the harm-minimizing level. For example, let S_t^a be the actual stock in period t, and write the flow damage as $aS_t^a + \frac{b}{2}S_t^{a2}$. Then $S^m \equiv -\frac{a}{b}$ is the cost-minimizing stock. Defining $S_t \equiv S_t^a - S^m$ we write damages as $\frac{b}{2}S_t^2$.

[15] Even in a richer model with gradual diffusion, the most important characteristic of the diffusion process for the policy ranking would be the amount of technology adopted during the current policy period relative to the long-term future. It is precisely this characteristic of technology diffusion that we capture by α. This model does not capture "endogenous" diffusion, where the tax level or the quota price effects the speed of diffusion.

[16] fn ARMAWe have $\hat{\theta}_t = \rho\ (\rho\theta_{t-2} + \alpha\varepsilon_{t-1} + (1-\alpha)\varepsilon_{t-1}) + \alpha\varepsilon_t = \rho\hat{\theta}_{t-1} + \rho(1-\alpha)\varepsilon_{t-1} + \alpha\varepsilon_t$.

in a larger future adoption.[17] Higher serial correlation of adopted technology implies that a cost shock today has a stronger impact on both BAU and optimal future levels of emissions. A shock that reduces future abatement costs lowers future emissions, thereby lowering future carbon stocks. That reduction lowers marginal damages, and causes a downwards shift of the SCC in Fig. 1. That shift partially offsets the welfare loss resulting from a cost shock under a quota. A shock that increases abatement costs similarly shifts the SCC up, again partially offsetting the welfare loss under a quota. Therefore, a larger ρ and a smaller α favor quotas.

The firms' emission benefits are $(h_t + \hat{\theta}_t)E_t - \frac{f}{2}E_t^2$, where f is the slope of the marginal abatement cost curve and $h_t + \theta_t$ is the technology stock; h_t is an exogenous deterministic trend and θ_t is the stochastic component. A higher value of $\hat{\theta}_t$ corresponds to a larger marginal benefit from emitting: a larger marginal abatement cost. A better-than-expected technological innovation therefore corresponds to a *negative* realization of the shock ε.

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 $\mathbb{E}_t\alpha\varepsilon_t = 0$)

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

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 $h_t + \hat{\theta}_t - fE_t = \tau_t$. This first order condition results in the firm's decision rule

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

It is convenient to model the tax-setting regulator as choosing *expected emissions* e_t^T, which is equivalent to setting the tax $\tau_t = h_t + \rho\theta_{t-1} - fe_t^T$. 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[18]

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

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. She balances the persistent costs from pollution with the transitory benefits from emitting. The discount factor is β. At the end of period t the regulator learns the value of θ_t by observing the permit price induced by the quota or the level of emissions induced by the tax. Thus, the regulator knows θ_t when choosing the policy level at $t+1$. The pollution stock is public information.

3.2. Policy ranking

We define the Social Cost of Carbon (SCC) for the dynamic model and show that its slope is much greater than the slope of the flow marginal damage. We then present and discuss the ranking criterion, which depends on both the slope ratio and the shift. The SCC in the climate setting, is linear in the state variables, i.e. it is of the form

$$SCC_t = \beta\mathbb{E}\left(\chi_{t+1} + \lambda S_{t+1} + \mu\theta_t\right),$$

Eqs. (12) and (13) in Appendix B provide the formulae (endogenous functions of model primitives) 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 realization. The appendix shows that both are positive constants. The parameter λ in the dynamic setting corresponds to g, the slope of the MD curve in the lower panels of Fig. 1. The parameter μ in the dynamic setting corresponds to ψ in the static setting. We took g and ψ as exogenous, whereas here we recognize that λ and μ are endogenous. Eq. (14) in Appendix B provides the formula for the intercept of the SCC, χ_t. We need this expression to calculate the optimal tax, but not for the policy ranking. 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 observers 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 *flow* 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 $r = 5.2\ 10^{-5}$ for our baseline calibration (Section 4). 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.[19] Lemma 1 gives the relation between these two slopes.

---

17 Appendix C.7 shows that $corr(\hat{\theta}_t, \hat{\theta}_{t+j}) = \rho^j\left(\rho^2\frac{\rho^{2j}-1}{\rho^2-1} + \alpha^2\right)^{-0.5}$ and confirms that this function decreases in α and increases in ρ.

18 We obtain this expression by replacing E with E^T in the firm's payoff. Using the definition of E^T, taking expectations, and then subtracting $\frac{b}{2}S_t^2$ (which is independent of current emissions), gives the flow payoff.

19 It is instructive to consider the time step and units explicitly when defining R. Appendix C.1 uses a parameter ϕ to denote the time step, enabling a simple scaling of the period's length. There, we define $R \equiv \frac{\lambda}{f}\phi$. R relates the slope of the SCC curve, $\frac{\partial SCC}{\partial S_t}$, to the slope of the MAC curve, $\frac{\partial MAC}{\partial E_t\phi}$. Here, $E_t\phi$ is the amount of emissions over the course of the period, equal to the annual emissions flow times the number of years in a period: we have to compare the cost of the marginal unit change of atmospheric carbon with a unit change of abatement over the course of the period (rather than with the annual *flow*). The parameter ϕ carries the unit time, and $R \equiv \frac{\partial SCC}{\partial S_t} / \frac{\partial MAC}{\partial E_t\phi} = \frac{\lambda}{f}\phi$ is unit free. In the main text, we set $\phi = 1$ (rather than "1 year") for ease of notation. This choice picks units in which years are normalized to unity.

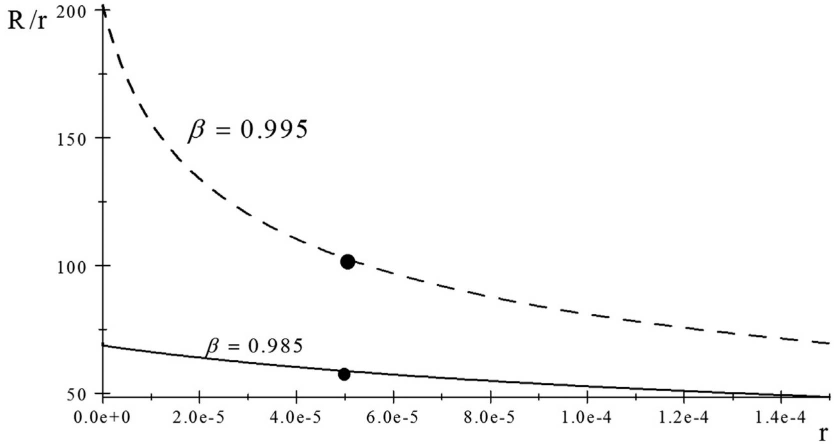

Fig. 2. The ratio of $\frac{R}{r}$ for $\delta = \rho = 1$, a one-year time step, and two annual discount factors. The heavy dots identify our baseline value $r = \frac{b}{f} = 5.2\ 10^{-5}$.

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*

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

Unsurprisingly, the relation between the flow ratio r and the stock ratio R depends on the discount factor β and the persistence of the pollutant δ. Fig. 2 graphs $\frac{R}{r}$ as a function of the flow pollution ratio r, using a one-year time step (the length of a period) and two alternative annual discount factors. (We set $\delta = 1$, the value corresponding to the TCRE model described in Section 4.) The heavy dots in Fig. 2 show $\frac{R}{r}$ at $r = 5.2\ 10^{-5}$, our baseline value, for a 1.5% and a 0.5% annual discount rate. Aggregate damages are *more convex* than flow damages: the SCC is much steeper in emissions than is the flow marginal damage curve.[20]

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

Proposition 1. *Taxes dominate quotas if and only if*

$$R < \frac{1}{\beta} - \frac{2\mu}{\alpha} \qquad \Leftrightarrow \qquad R < R^{crit} \equiv -\frac{1}{2}\kappa_1 + \frac{1}{2}\sqrt{\kappa_1^2 + 4\kappa_0}$$
$$\text{with } \kappa_1 \equiv \frac{\delta\rho(2-\alpha)}{\alpha} \text{ and } \kappa_0 \equiv \frac{1-\beta\delta\rho}{\beta^2}. \tag{3}$$

For flow pollutants, taxes dominate quotas if and only if $r < 1$. The first condition in Proposition 1 shows that: (i) the relevant slope in the ranking criterion for a stock pollutant is R instead of r; (ii) a higher responsiveness μ of the SCC to the technology shock favors quotas; and (iii) slow technology adoption (small α) favors quotas.[21] The endogenous μ is the shadow value of the interaction term $\theta_t S_t$. It measures the responsiveness of the social cost of marginal emissions (SCC) to the technology realization. The right hand side of the equivalence (3) expresses the ranking 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.

Fig. 3 graphs R^{crit} as a function of the "joint persistence", $\delta\rho$, of the stock pollutant and technology for three values of α, with an annual time step and the discount factor $\beta = 0.985$. With little or no pollution or technology persistence, $\delta\rho \approx 0$ 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 Fig. 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, but even if R is close to 0. Section 4 further explores this possibility.

We provide intuition for our results using the case of a *flow pollutant*, where 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; here the emission quantity overreacts to a cost shock, 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 the magnitude of the deadweight loss. Under quotas, emissions are constant, but the firms' equilibrium

[20] With a flow pollutant, E, the slope of the marginal damage of an additional unit of emissions is b. With a stock pollutant, S, the marginal damage of an additional unit of emissions is $\lambda \frac{\partial S_{t+1}}{\partial E_t} = \lambda$.

[21] Eq. (13) in the appendix provides the formula for μ in terms of the model's fundamentals. Importantly, μ and λ are independent of α.

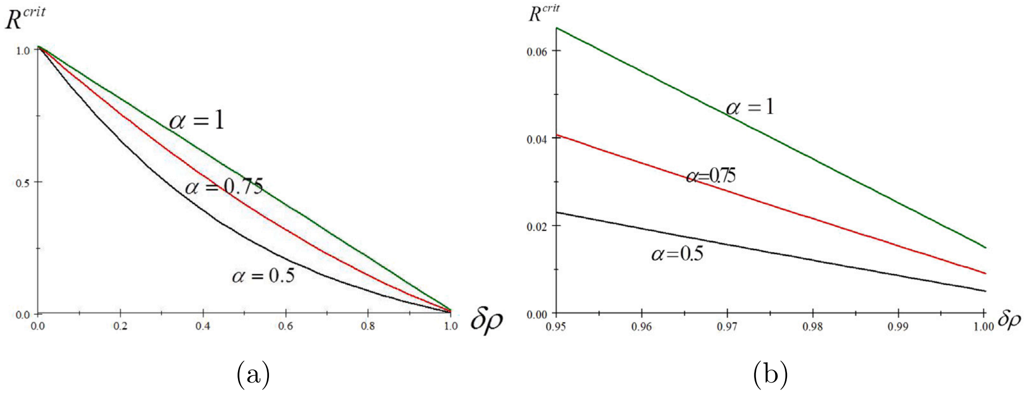

Fig. 3. Panel (3(a)) graphs R^{crit} as a function of $\delta\rho$ for three values of α, using a one-year time step and an annual discount rate of 1.5%. Panel (3(b)) shows the same graph for values of $\delta\rho$ close to 1, the relevant case for climate change.

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 the damage convexity is greater for the *stock pollution* than for the flow pollution ($R \gg r$).

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) *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, a persistent technological innovation today implies lower future emissions, resulting in a lower future pollution stock.[22] Consequently, a technological innovation *reduces the marginal damages* (the SCC) resulting from an additional emission unit today. 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 shifts in marginal costs and marginal damages resulting from the innovation are perfectly correlated, differing only in magnitude. A part of what would be an "overreaction" of emission prices under quantity regulation of a flow pollutant becomes a socially optimal variation under a stock pollutant.

Proposition 1 shows that a higher value of μ favors quotas. The endogenous value μ measures the responsiveness of the social cost of marginal emissions to technology realization. It is the derivative of the SCC with respect 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 equilibrium marginal 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 Fig. 1), it amplifies the optimal price fluctuations in response to the innovation, relative to the case of the flow pollutant (upper panels of Fig. 1). Indeed, for $\alpha = 1$, the left side of the policy-ranking equivalence (3) (dynamic model) reproduces the left side of the graph-based equivalence (1) 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 stock pollution extension in Section 2.2 and the dynamic model is that both R and μ are endogenous in Eq. (3), whereas Section 2.2 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 extension in Section 2.2 cannot capture the fact that technology diffusion takes more than one period ($\alpha < 1$).

Before continuing the discussion of technology diffusion and 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 $\alpha < \alpha^{crit}$.

[22] In 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

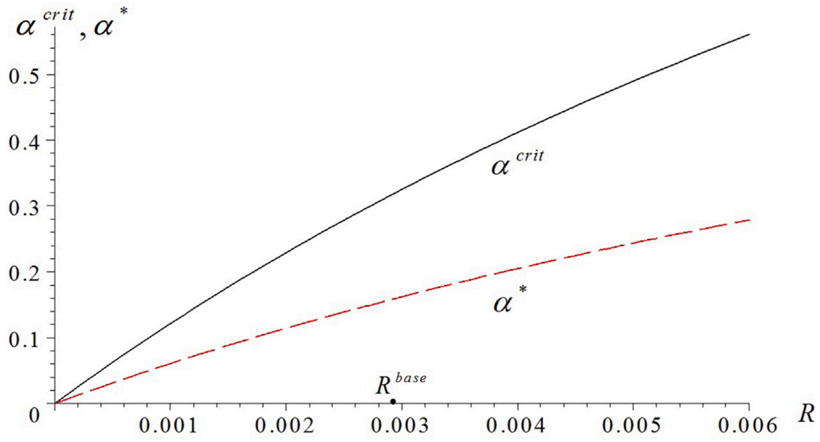

Fig. 4. Quotas dominate taxes for $\alpha < \alpha^{crit}$, the solid graph. The quota is first best for $\alpha = \alpha^*$, the red dashed graph. The graphs use an annual time step and baseline values $\beta = 0.985$, $\delta = 1$. R^{base} marks the slope ratio in our baseline calibration.

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 ($\alpha < 1$). Fig. 4 graphs α^* as a function of R, the ratio of stock damage convexity to abatement cost convexity. It also graphs the critical diffusion level α^{crit}, below which quotas dominate taxes,

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. We noted above (and in Footnote 17) that partial diffusion increases the correlation between current and future adopted technology. As a result, a given level of adoption today signals even more future adoption. The socially optimal level of 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 C.6 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.

3.3. A welfare measure

Although the expected emissions trajectories are the same under optimal taxes and quotas, asymmetric information causes welfare to differ under the two policies. The difference in welfare is proportional to σ^2, the variance of the cost shock. Lacking a good estimate of σ^2, we instead provide a measure of the *relative* welfare gain that is independent of σ^2.[23]

To this end, we consider the first-best (full information) setting, where the regulator learns the current shock at the same time as firms do, before choosing the current policy; there, the tax and the quota are equivalent. This problem is still stochastic, because the regulator knows only the expectation of future emissions, not their realization. However, full information eliminates asymmetric information between firms and the regulator.

We define the relative welfare gain in moving from quotas to taxes as a ratio: the loss due to moving from full to asymmetric information conditional on using quotas, relative to the loss due to moving from full to asymmetric information conditional on using taxes:

$$G(\alpha) \equiv \frac{V\,(\text{full info}) - V\,(\text{asym. info, quota})}{V\,(\text{full info}) - V\,(\text{asym. info, tax})}, \tag{4}$$

where $V(\cdot)$ denotes the present discounted stream of benefits, conditional on the policy scenario. The function G is non-negative, and it is greater than 1 if and only if taxes dominate quotas.

[23] We know of only two studies, both over a decade old, that use values of σ^2 to quantify welfare changes in the linear-quadratic setting. However, the models in those papers differ so much from ours that we cannot sensibly adapt their estimates to our model.

Proposition 3. *The relative welfare gain of using taxes instead of quotas is*

$$G = \left(\frac{(\alpha - \beta\mu)}{\beta(\alpha\lambda + f\mu)} f \right)^2 . \tag{5}$$

Recalling that the endogenous variables λ and μ are independent of α, Eq. (5) shows that the relative welfare gain, $G(\alpha)$, is decreasing for $\alpha < \alpha^*$ and increasing for $\alpha > \alpha^*$; it reaches its minimum at $\alpha = \alpha^*$, where $G(\alpha^*) = 0$. There, the quota is first best: asymmetric information then creates no loss under quotas, although there remains a loss under taxes. Proposition 2.ii identifies α^{crit} as the value of $\alpha > \alpha^* > 0$ at which the payoff is the same under taxes and quotas: $G(\alpha^{crit}) = 1$.

As technology diffusion α becomes small, current innovations ϵ_t have only a small impact on the current period's marginal abatement cost. In Fig. 1, the MAC curve shifts less and less as $\alpha \to 0$. As the MAC curve stops shifting, the policymaker sets the emissions allocation without any uncertainty. Therefore, the welfare loss under a tax approaches the same welfare loss as under a quota and $G \to 1$ as $\alpha \to 0$.[24]

4. The climate application

This section quantifies our results above, using the calibration in Karp and Traeger (2024). Costs and damages are based on Nordhaus and Sztorc (2013), and climate dynamics are based on the somewhat better-performing TCRE model (IPCC, 2013; Dietz and Venmans, 2018). Appendix A graphs the temperature impulse response of this TCRE-based climate model and compares it to a middle-of-the road climate model as well as the 2016 DICE model. Here, we give a brief summary of the calibration; C.5 provides details.

The "transient climate response to cumulative carbon emissions" (TCRE) makes global warming a linear function of the cumulative past emissions (*not* the carbon concentration in the atmosphere). Emissions gradually leave the atmosphere, but they have a cumulative effect on temperature. These two non-linear effects almost cancel each other, resulting in an almost-linear relation between cumulative emissions and the temperature. With the TCRE model, we (i) track cumulative historic emissions so that $\delta = 1$ and (ii) associate the stock variable S_t with temperature. Our baseline uses the IPCC's (2021) best guess transient climate response to cumulative carbon emissions of $TCRE = 1.65 \, \frac{^{\circ}\mathrm{C}}{\mathrm{TtC}}$; this value coincides with the midpoint of the IPCC's (2013) interval estimate of the transient climate response of $[0.8, 2.5]$ in $\frac{^{\circ}\mathrm{C}}{\mathrm{TtC}}$, which we will relate to for our "greater climate response" scenario, and coincides with the point estimate of the IPCC's (2021).

We use a five-year policy period and assume an annual discount rate of 1.5%. We set $\rho = 1$ because our model describes the role of technological progress that persistently alters abatement costs. Our other baseline calibration follows DICE in assuming that flow damages are zero at the pre-industrial temperature ($T = 0$) and that damages at $T = 2$ equal approximately 1% of world output. Global world output is 130 trillion USD in 2020. Our baseline also adopts Karp and Traeger's (2024) estimate of the DICE's abatement cost slope as $f = 2.5\,10^{-9} \, \frac{\mathrm{USD}}{\mathrm{tCO}_2^2}$. We require the estimate of the technology level $h_{2020} = 101 \, \frac{\mathrm{USD}}{\mathrm{tCO}_2}$ only to calculate the social costs of carbon for our calibration. The calibration assumes that this intercept falls exogenously by 1% per year.

We also consider three alternatives.[25] The scenario "greater climate response" (greater CR) uses all of the baseline assumptions except that it sets the TCRE to the upper bound of the IPCC's (2013) estimated range, $TCRE = 2.5$. The scenario "greater damage convexity" (greater DC) uses the baseline parameters except that it assumes that a 1 degree temperature anomaly creates zero damage, but a three degree anomaly creates damages equal to 5% of world output. The final scenario assumes both greater damage convexity and greater climate response, combining the changes of the second and third scenarios.[26]

Karp and Traeger (2024) regress carbon emissions on green patents, producing an estimate of $\alpha \approx 0.3$ for a five-year time step. This estimate is consistent with the widely accepted view that technology diffuses with a lag. Here we compare taxes and quotas for a range of α.

Our calibration implies 2020 BAU emissions of 40 GtCO_2, slightly higher than estimated emissions, and thus consistent with the current weak climate policy. Optimal 2020 emissions range from 22–29 GtCO_2, over the four scenarios, implying reductions of 27–45% relative to BAU. The optimal taxes range from 27–43 $\frac{\mathrm{USD}}{\mathrm{tCO}_2}$. These values are independent of α.

Table 1 identifies the optimal policy instrument, Quota or Tax, under the four scenarios, for three values of α. The row α^{crit} shows the critical value of α, below which quotas dominate taxes. In the baseline, taxes dominate quotas if adoption happens reasonably fast. In the scenarios with greater climate response or greater damage convexity, quotas dominate taxes even if only half of the innovation is adopted within the regulation period. In the scenario with both greater climate response and greater damage convexity, quotas always dominate taxes. The row α^* shows the value of α at which the quota produces the full information first-best outcome.

24 Even with α approaching zero, the welfare costs of asymmetric information remain positive. Current innovations still translate into future emission reductions and, thus, affect future marginal damages and today's SCC. In Fig. 1, the SCC curve still shifts even if the current MAC curve does not (for $\alpha \to 0$). Under full information, the policymaker would already know the shift of the SCC curve at the point of deciding the targeted allocation. In our setting with $\alpha \to 0$, the policy maker misses this information, and welfare is less than under full information.

25 The scenario "greater damage convexity" corresponds to the "concerned" scenario in Karp and Traeger (2024). The other two scenarios are newly calibrated for the current study.

26 The value of b varies across the scenarios. We have $b = 1.32\,10^{-13}$ (baseline); $b = 3.02\,10^{-13}$ (greater climate response); $b = 6.58\,10^{-13}$ (greater damage convexity); and $b = 1.51\,10^{-12}$ (both greater climate response and damage convexity).

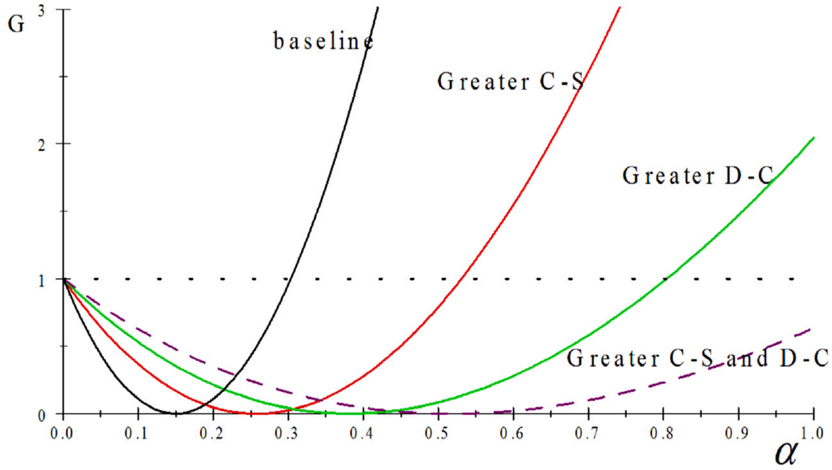

Fig. 5. $G(\alpha)$ equals the ratio of the gain in moving from quotas to full information, relative to the gain in moving from taxes to full information. Quotas dominate taxes if and only if $G < 1$, as occurs for $\alpha < \alpha^{crit}$. The curves reach their minimum at $\alpha = \alpha^*$, where $G = 0$; at this value of α, the quota is first-best.

Table 1
The table presents the optimal policy choice, Quota or Tax, for two different pure rates of time preference (prtp) and three different choices for the share of technological innovation adopted within the regulation period ($\alpha = 1, 0.5, 0.25$). The table also presents α^{crit}, below which quotas dominate taxes, and α^* where a quota produces the full-information first best outcome (no deadweight loss).

	Baseline	Greater climate response	Greater damage convexity	Greater climate response & Damage convexity
p.r.t.p. 1.5%				
$\alpha = 1$:	Tax	Tax	Tax	Quota
$\alpha = 0.5$	Tax	Quota	Quota	Quota
$\alpha = 0.25$	Quota	Quota	Quota	Quota
α^{crit}	0.3	0.53	0.8	>1
α^*	0.15	0.26	0.38	0.52
SCC in USD/tCO_2	27	40	29	42
p.r.t.p. 0.5%				
for any $\alpha \leq 1$	Quota	Quota	Quota	Quota
SCC in USD/tCO_2	55	69	49	60

The lower section of the table reduces the pure rate of time preference from 1.5% to 0.5%, following the median response of Drupp et al.'s (2018) survey. Under such an increased attention to future damages, quotas always dominate taxes. The volatility of emissions under a tax gains in weight relative to the volatility of firms' abatement costs under a quota. Finally, Table 1 presents the expected optimal carbon price in the different scenarios. Our baseline's 2020 price of $27\frac{\text{USD}}{\text{tCO}_2}$ equals Dietz et al.'s (2021) price of $27\frac{\text{USD}}{\text{tCO}_2}$ deriving from an enhanced version of the DICE model excluding certain non-linear temperature-carbon cycle feedbacks and lies slightly below the average 2020 price of the EU ETS of $30\frac{\text{USD}}{\text{tCO}_2}$ (World Bank, 2021). The other scenarios yield higher expected SCCs. Maybe surprisingly, the greater damage convexity reduces the SCC from $55\frac{\text{USD}}{\text{tCO}_2}$ to $49\frac{\text{USD}}{\text{tCO}_2}$ under a low pure rate of time preference. The scenario with higher damage convexity reduces damages at low temperature and increases damages at high temperatures. Under the reduced time preference, optimally regulated temperatures are sufficiently low such that the expected SCC falls relative to the baseline.

Fig. 5 provides another perspective, showing graphs of the function G, defined in Proposition 3. Quotas dominate taxes if and only if $G < 1$. The graphs intersect the horizontal line $G = 1$ (where the two policies produce the same level of welfare) at $\alpha = \alpha^{crit}$ and they reach their minimum $G = 0$ (where there is no deadweight loss under quotas) at $\alpha = \alpha^*$.

Compared to previous results, our estimates are more favorable to quotas for three reasons. First, our model of gradual adoption of technology favors quotas when $\alpha < 1$. The gradual adoption of technology and the resulting gradual revelation of otherwise-hidden information reduces one of the major disadvantages of quotas: the concern that aggregate emissions respond too slowly to firms' information.

Second, the TCRE model implies that the state variable does not decay: $\delta = 1$. Simplified models merely relying on the atmospheric carbon stock as the state variable have lower persistence factors, reflecting the removal of the atmospheric carbon. The temperature response even in Nordhaus's (2017) DICE model is too sluggish, performing worse than the TCRE model's immediate persistent response (Mattauch et al., 2020; Dietz et al., 2021). Policy ranking can be sensitive to small changes in δ and β, parameters that determine the future costs of current actions.

Third, we emphasize technology rather than the other more transitory shocks that affect firms' emissions decisions. Thus, we set $\rho = 1$; ρ would be smaller if the shock were an amalgam of both persistent and transitory shocks. Again, our rationale for this focus is that the policy can be conditioned on the more transitory shocks, because these are publicly observed when policy is implemented, although not when policy rules are chosen.

5. Conclusions

A widely used (static) criterion for ranking price-based and quantity-based regulation does not carry over to the dynamic setting where current shocks affect future abatement costs, thereby affecting future regulation. We considered a setting with asymmetric information between the regulator and firms arising from technological change. The policy maker regulates an externality but does not observe current innovations. The standard ranking criterion incorporates the effect of innovations on firms' cost structure. Our criterion recognizes that the current technology innovation also alters future abatement costs and abatement levels, changing the stock trajectory. Both the persistent impact of shocks and a delayed technology diffusion favor quantity regulation.

Our discussion focuses on pollution control to mitigate climate change, where Weitzman's (1974) static ranking criterion is often informally applied. 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 a simple criterion for ranking prices versus quantities for stock externalities under asymmetric information. This criterion depends on both the ratio of the slopes and on the ratio of the shock-induced shifts in the intercepts of the marginal damage and abatement cost curves. The ratio of slopes is a familiar component, but the ratio of shifts in intercepts is novel, and is equally important in determining the ranking. 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. These conceptual changes in the ranking criterion result from the persistence of technology and its (potentially) gradual diffusion.

Our empirical application shows that the conceptual correction of the ranking criterion 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 criterion has serious policy implications.

Technological uncertainty, which lies at the heart of Weitzman's (1974) asymmetric information problem, 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.

CRediT authorship contribution statement

Larry Karp: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Project administration, Resources, Validation, Visualization, Writing – original draft. **Christian Traeger:** Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Resources, Validation, Visualization, Writing – original draft, Writing – review & editing.

Appendix A. Climate dynamics

This section discusses our model's implied climate dynamics. Fig. 6 shows the temperature response over the coming 150 years to 100Gt of carbon dioxide emitted today. The experiment fixes the background concentration at today's atmospheric carbon dioxide concentration.[27]

Our base scenario, taken from Karp and Traeger (2024), matches the IPCC's (2021) best guess for the transient climate response to cumulative carbon emissions, a 1.65C warming for a 1000Gt emission pulse and, thus, one tenth of this response to a 100 Gt impulse. The thick green line labeled "Bern 2.5 PD" presents the results of the Bern 2.5 model, a detailed middle-of-the-road climate change

[27] The TCRE model and the DICE model's response is independent of the background concentration, which matters for the Bern 2.5 model where "PD" abbreviates "present day". We are grateful to Fortunat Joos who provided the Bern 2.5 results. Similar data is presented in Traeger (2023) and Karp and Traeger (2024)

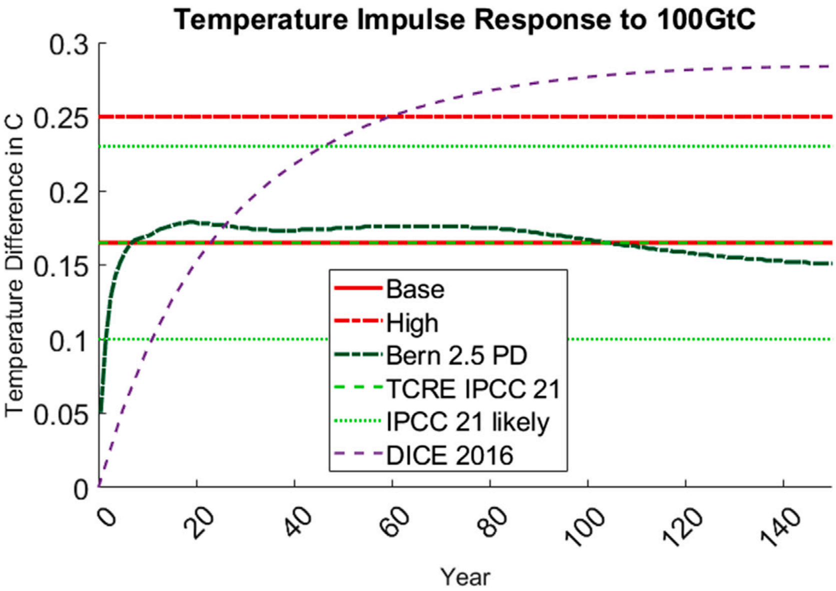

Fig. 6. The figure shows the temperatures impulse response to carbon dioxide emissions. Our base scenario coincides with the IPCC's (2021) best guess for the transient climate response to cumulative carbon emissions, which is a good approximation to more sophisticate climate change models like the Bern 2.5 model. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

model used in previous iterations of the IPCC reports to calculate global warming potentials. It shows that the actual temperature response is not expected to be a truly flat line, but that the TCRE's constant temperature response is a good approximation to sophisticated climate change models. The graph also depicts the IPCC's (2021) 67% likelihood interval (green dotted). Our "greater climate response" scenario (high, red dash-dotted) lies just outside of this interval, representing the upper bound of the IPCC's (2013) interval estimate for the TCRE. Finally, the graph compares these models to the impulse response of Nordhaus's (2017) DICE 2016 model. As emphasized by previous authors (Mattauch et al., 2020; Dietz et al., 2021), the DICE model's response to carbon emissions is too sluggish, and the medium run temperature response of the DICE 2016 model overshoots the best guess even more than our "high" scenario.

Appendix B. Derivations and proofs

Derivation of Eq. (1):

Using the upper right panel of Fig. 1, it is straightforward 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.

We use the lower right panel in Fig. 1 to confirm inequality (1). The ranking criterion depends on both the responsiveness φ of the MD curve to a shift of the MAC and on the relative slopes. 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

$$d = \frac{\varphi\, k}{1-\varphi} \quad \text{or} \quad \frac{d}{k} = \frac{\varphi}{1-\varphi}. \tag{6}$$

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 Eq. (6) to replace the fraction $\frac{d}{k}$ delivers

$$\frac{h\, m^{MAC}}{k} = \frac{1}{1-\varphi}. \tag{7}$$

Similarly, we observe that $h\, m^{MD} = s - d \Rightarrow \frac{h\, m^{MD}}{s} = 1 - \frac{d}{s}$. Using Eq. (6) to replace d, we obtain

$$\frac{h\, m^{MD}}{s} = 1 - \frac{\varphi}{1-\varphi}\frac{k}{s}. \tag{8}$$

Dividing Eq. (8) by Eq. (7) 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., $\frac{s}{k} < 1$, leading to Eq. (1).

Notation for proofs of dynamic model:

We take advantage of the linear-quadratic structure to unify the separate the problems when the regulator uses taxes or quotas or in the first best (full information) setting. To this end, we introduce the indicator function $\Phi = 1$ under a tax and $\Phi = 0$ under a quota. Using $x_t \in \{e^T, e^Q\}$ to denote the regulator's control under tax and quantity regulation, respectively, the regulator's problem, for $i \in \{T, Q\}$, is

$$\max \mathbf{E}_t \sum_{\tau=0}^{\infty} \beta \left[\left(h_{t+\tau} + \rho\theta_{t+\tau-1}\right) x_{t+\tau} - \frac{f}{2}\left(x_{t+\tau}\right) + \Phi\frac{\alpha^2}{2f}\sigma^2 - \frac{b}{2}S_{t+\tau}^2\right]\phi$$

$$\text{subject to } S_{t+\tau+1} = \delta S_{t+\tau} + \phi x_{t+\tau} + \Phi\phi\alpha\frac{\varepsilon_t}{f} \text{ and } \theta_t = \rho\theta_{t-1} + \varepsilon_t.$$

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. The problem formulated using x and Φ is 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 $Y_t = \left(S_t, \theta_{t-1}\right)'$ and we define:

$$Q = \begin{pmatrix} -b & 0 \\ 0 & 0 \end{pmatrix}, \quad A = \begin{pmatrix} \delta & 0 \\ 0 & \rho \end{pmatrix}, \quad W = \begin{pmatrix} 0 & \rho \end{pmatrix},$$
$$B = \begin{pmatrix} \phi \\ 0 \end{pmatrix}, \quad C = \begin{pmatrix} \Phi\phi\frac{\alpha}{f} \\ 1 \end{pmatrix}. \tag{9}$$

The net flow payoff and equation of motion for the generic problem are:

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

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

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

$$J_t^i\left(Y_t\right) = \text{Max}_{x_t}\left[h_t x_t - \frac{1}{2}f x_t^2 + \frac{1}{2}Y_t'QY_t + WY_t x_t + \Phi\frac{\alpha^2}{2f}\sigma^2\right]\phi + \beta\mathbf{E}_t J_{t+1}^i\left(Y_{t+1}\right). \tag{10}$$

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

The value function for the LQ problem, for $i \in \{T\,Q\}$, is linear-quadratic: $J_t^i\left(Y_t\right) = V_{0,t}^i + V_{1t}'Y_t + \frac{1}{2}Y_t'V_2Y_t$. 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_{0t}^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:

$$V_2 = -\begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix}. \tag{11}$$

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

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

Online Appendix C.2 provides the details of the following steps:

(1) We substitute the equations of motion into the right side of the DPE, Eq. (10), 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 Riccati 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

$$\lambda = \frac{1}{2\beta\phi}\left(-\varpi + \sqrt{\varpi^2 + 4\beta\phi^2 bf}\right) > 0 \tag{12}$$

$$\mu = \frac{\lambda}{f}\left(\frac{\phi\beta\delta\rho}{(1-\beta\delta\rho) + \beta\phi\frac{\lambda}{f}}\right). \tag{13}$$

From inspection of Eq. (12), $\lambda > 0$, so the numerator of the right side of Eq. (13) is positive. Therefore, μ has the same sign as ρ, which in our setting is positive, because the shock describes a technological innovation.[28].

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 Eq. (12) by f establishes Lemma 1. ■

The algorithm described above also produces the formula for χ_t, the intercept of the SCC_t (Online Appendix C.2) Using the definitions

$$M = \frac{\beta\delta f}{f + \beta\phi\lambda} \text{ and } N = -\frac{\beta\delta\phi\lambda}{f + \beta\phi\lambda}.$$

we express χ_t as

$$\chi_t = -N \sum_{j=0}^{\infty} M^j h_{t+j}. \tag{14}$$

If h falls at a constant rate (as in our climate application) we can express this infinite sum as a function of the current level h_t and the model parameters.

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

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

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

(Online Appendix C.2 derives this relation; see equation 29.) 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

$$\Delta_t = V_{0,t}^T - V_{0,t}^Q = \tfrac{\alpha^2}{2f}\sigma^2\phi + \beta\Delta_{t+1}$$

$$-\tfrac{1}{2}\beta\sigma^2\left[\begin{pmatrix} \phi\frac{\alpha}{f} & 1 \end{pmatrix}\begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix}\begin{pmatrix} \phi\frac{\alpha}{f} \\ 1 \end{pmatrix} - \begin{pmatrix} 0 & 1 \end{pmatrix}\begin{bmatrix} \lambda & \mu \\ \mu & \nu \end{bmatrix}\begin{pmatrix} 0 \\ 1 \end{pmatrix}\right] \Rightarrow$$

$$\Delta_t = \beta\Delta_{t+1} + \tfrac{\alpha^2}{2f}\sigma^2\phi - \tfrac{1}{2}\beta\sigma^2\phi\alpha\tfrac{\phi\alpha\lambda + 2\mu f}{f^2} = \beta\Delta_{t+1} + \tfrac{\alpha\phi}{2f}\sigma^2\left(\alpha - \beta\tfrac{\phi\alpha\lambda + 2\mu f}{f}\right).$$

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

$$\Delta = \frac{1}{1-\beta}\frac{\alpha\phi}{2f}\sigma^2\left(\alpha - \beta\frac{\phi\alpha\lambda + 2\mu f}{f}\right). \tag{15}$$

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

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

This equation implies that taxes dominate quotas if and only if

$$\alpha - \beta\left(2\mu + \alpha R\right) > 0. \tag{16}$$

Rearranging this inequality establishes the equivalence (3). The explicit dependence on the time step ϕ has dropped out; ϕ matters only through the scaling of the discount factor β and the decay factor δ.

Using inequality (16), the definition of R, and Eq. (13) we have

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

Multiplying by the positive denominator, this inequality is equivalent to

[28] The units of λ are $\frac{\text{USD}}{\text{GtCO}_2{}^2}$: 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.

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

where we use 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} in the proposition. Hence the inequality is satisfied for $R\in[0,R^{crit})$.

Proof of 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 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:

$$\rho\theta_{t-1}+\alpha\varepsilon_t-fE_t^{FB}=\beta\left(\chi_t+\lambda\left(\delta S_t+E_t^{FB}\right)+\mu\left(\rho\theta_{t-1}+\varepsilon_t\right)\right), \tag{17}$$

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 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 Eq. (17). The fact that χ_t, λ, and μ are the same in the first best world and under quotas (and 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 Eq. (17) 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

$$\beta\mu\leq 1\Leftrightarrow\beta^2 R\frac{\delta\rho}{(1-\beta\delta\rho)+\beta R}\leq 1\Leftrightarrow$$
$$R\beta\left(\beta\delta\rho-1\right)\leq(1-\beta\delta\rho). \tag{18}$$

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

$$\frac{dR^{crit}}{d\alpha}=-\frac{1}{2}\frac{\kappa_1-\sqrt{\kappa_1^2+4\kappa_0}}{\sqrt{\kappa_1^2+4\kappa_0}}2\delta\frac{\rho}{\alpha^2}>0. \tag{19}$$

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

$$\frac{dE^T}{d\varepsilon_t}=\frac{\alpha}{f}>\frac{\alpha-\beta\mu}{f+\beta\lambda}=\frac{dE^{FB}}{d\varepsilon_t}, \tag{20}$$

where the second equality uses the first order condition (17) 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 (19) 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 $\alpha<\alpha^{crit}$. ■

Proof. Proposition 3 (Sketch) Online Appendix C.3 Provides the details. Calculation of the expected welfare under full information requires solving a standard linear-quadratic control problem, in which the regulator learns the current shock, ε_t, at the same time as firms. Comparing this payoff to the expected level of welfare under a quota (which we solved for the proof of Proposition 1) gives the numerator of the ratio G, shown in Eq. (4). Subtracting the payoff differences under taxes versus quotas (which we obtained from the proof of Proposition 1) from the numerator produces the denominator of G. Simplifying this ratio produces Eq. (5).

Appendix C. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jeem.2024.102951.

References

Aldy, J.E., Krupnick, A.J., Newell, R.G., Parry, I.W., Pizer, W.A., 2010. Designing climate mitigation policy. J. Econ. Lit. 48 (4), 903–934.
Anderson, E., Brock, W., Hansen, L.P., Sanstad, A.H., 2014. Robust analytical and computational exploration of coupled economic-climate models with carbon-climate response. RDCEP Working Paper No.13-05.
Boleslavsky, R., Kelly, D., 2014. Dynamic regulation design without permits: the importance of timing. J. Public Econ. (120), 169–180.
Brock, W., Xepapadeas, A., 2017. Climate change policy under polar amplification. Eur. Econ. Rev. 99, 93–112, Combating Climate Change. Lessons from Macroeconomics, Political Economy and Public Finance.
Burtraw, D., Holt, C., Palmer, K., Shobe, W., 2020. Quantities and prices: Price-responsive allocance supply in environmental markets. RFF Working Paper 20-17.
Dietz, S., van der Ploeg, F., Rezai, A., Venmans, F., 2021. Are economists getting climate dynamics right and does it matter? J. Assoc. Environ. Resour. Econ. 8 (5).
Dietz, S., Venmans, F., 2018. Cumulative carbon emissions and economic policy: in search of general principles. Grantham Working Paper 283.
Doda, B., 2016. How to price carbon in good times... and bad! Clim. Chang. 2016-01 7 (1), 135–144.
Drupp, M., Freeman, M., Groom, B., Nesje, F., 2018. Discounting disentangled. Am. Econ. J. Econ. Policy 4 (4), 109–134.
Ellerman, A.D., Wing, I.S., 2003. Absolute versus intensity-based emission caps. Clim. Policy 3, S7–S20.
Fischer, C., Springborn, M., 2011. Emissions targets and the real business cycle: Intensity targets versus caps or taxes. J. Environ. Econ. Manag. 62, 352–366.
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. Clim. Chang. Econ. 04 (03), 1350010, 1–28.
Hepburn, C., 2006. Regulation by prices, quantities, or both: a review of instrument choice. Oxf. Rev. Econ. Policy 22 (2), 226–247.
Heutel, G., 2012. How should environmental policy respond to business cycles? Optimal policy under persistent productivity shocks. Rev. Econ. Dyn. 15, 244–264.
Hoel, M., Karp, L., 2001. Taxes and quotas for a stock pollutant with multiplicative uncertainty. J. Public Econ. 82, 91–114.
Hoel, M., Karp, L., 2002. Taxes versus quotas for a stock pollutant. Resour. Energy Econ. 24, 367–384.
IPCC, 2013. Climate Change 2013 - The Physical Science Basis - Contribution of Working Group 1 to the IPCC - Summary for Policymakers. Technical Report, IPCC.
IPCC, 2021. In: Masson-Delmotte, V., Zhai, P., Pirani, A., Connors, S.L., Péan, C., Berger, S., Caud, N., Chen, Y., Goldfarb, L., Gomis, M.I., Huang, M., Leitzell, K., Lonnoy, E., Matthews, J., Maycock, T.K., Waterfield, T., Yelekçi, O., Yu, R., Zhou, B. (Eds.), Climate Change 2021: The Physical Science Basis. Cambridge University Press, Cambridge
Jotzo, F., Pezzey, J.C., 2007. Optimal intensity targets for greenhouse gas emissions trading under uncertainty. Environ. Resour. Econ. 38, 259–284.
Kalkuhl, M., Wenz, L., 2020. The impact of climate conditions on economic production. Evidence from a global panel of regions. J. Environ. Econ. Manag. 103, 102360.
Karp, L., Traeger, C., 2024. Smart cap. J. Eur. Econ. Assoc. (forthcoming).
Karp, L., Zhang, J., 2005. Regulation of stock externalities with correlated abatement costs. Environ. Resour. Econ. 32, 273–299.
Karp, L., Zhang, J., 2006. Regulation with anticipated learning about environmental damage. J. Environ. Econ. Manag. 51, 259–280.
Karp, L., Zhang, J., 2012. Taxes versus quantities for a stock pollutant with endogenous abatement costs and asymmetric information. Econom. Theory 49, 371–409.
Kelly, D.L., 2005. Price and quantity regulation in general equilibrium. J. Econom. Theory 125 (1), 36–60.
Keohane, N., 2009. Cap and trade, rehabilited: Using tradable permits to control U.S. greenhouse gasses. Rev. Environ. Econ. Policy 3, 42–62.
Mattauch, L., Matthews, H.D., Millar, R., Rezai, A., Solomon, S., Venmas, F., 2020. Comment on: Steering the climate system - Using inertia to lower the cost of policy. Amer. Econ. Rev. 114, 1231–1237.
Montero, J.P., 2002. Prices versus quantities with incomplete enforcement. J. Public Econ. 85, 435–454.
Montero, J.-P., 2008. A simple auction mechanism for the optimal allocation of the commons. Amer. Econ. Rev. 98 (1), 496–518.
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 externalitiues under uncertainty. J. Environ. Econ. Manag. 45, 416–432.
Newell, R.G., Pizer, W.A., 2008. Indexed regulation. J. Environ. Econ. Manag. 56, 221–233.
Nordhaus, W.D., 2008. A Question of Balance: Weighing the Options on Global Warming Policies. Yale University Press, New Haven, URL: http://nordhaus.econ.yale.edu/DICE2007.htm.
Nordhaus, W.D., 2017. Revisiting the social cost of carbon. Proc. Natl. Acad. Sci. 114 (7), 1518–1523.
Nordhaus, W., Sztorc, P., 2013. DICE2013R:Introduction and User's Manual. dicemodel.net.
Perino, G., Requate, T., 2012. Does more stringent environmental regulation induce or reduce technology adoption? J. Environ. Econ. Manag. 64 (3), 456–467.
Pizer, W.A., 1999. The optimal choice of climate change policy in the presence of uncertainty. Resour. Energy Econ. 21, 255–287.
Pizer, W., Prest, B., 2020. Prices versus quantities with policy updating. J. Assoc. Environ. Resour. Econ. 7, 483 518.
Requate, T., Unold, W., 2003. Environmental policy incentives to adopt advanced abatement technology: Will the true ranking please stand up? Eur. Econ. Rev. 47, 125–146.
Rogers, E.M., 2003. Diffusion of Innovations, fixth ed. Simon and Schuster.
Shinkuma, T., Sugeta, H., 2016. Taxes versus emissions trade scheme in the long run. J. Environ. Econ. Manag. 75, 12–24.
Stavins, R.N., 1996. Correlated uncertainty and policy instrument choice. J. Environ. Econ. Manag. 30 (2), 218–232.
Stavins, R.N., 2020. The future of U.S. carbon-pricing policy. Environ. Energy Policy Econ. 1, 8–64.
Traeger, C.P., 2023. ACE—Analytic climate economy. Am. Econ. J. Econ. Policy 15 (3), 372–406.
van der Ploeg, R., Emmerling, J., Groom, B., 2022. The Social Cost of Carbon with Intragenerational Inequality Under Economic Uncertainty. CESifo Working Paper.
Weitzman, M., 1974. Prices versus quantities. Rev. Econom. Stud. 41 (4), 477–491.
Weitzman, M., 2018. Prices or quantities dominate banking and borrowing. In: NBER 24218.
Wood, P.J., Jotzo, F., 2011. Price floors for emissions trading. Energy Policy 39, 1746–1753.
World Bank, 2021. Carbon price EU ETS 2020. data retrieved from the World Bank's Carbon Pricing Dashboard, https://carbonpricingdashboard.worldbank.org.
World Bank, 2023. Carbon pricing dashboard. 2022 data retrieved from the World Bank's Carbon Pricing Dashboard 03/08/2023, https://carbonpricingdashboard.worldbank.org.